

UTILITY PATENT APPLICATION TRANSMITTAL

Attorney Docket No.

663005.482C3

First Inventor or Application Identifier

Antony S. Williams

Title

A METHOD AND SYSTEM FOR REGISTERING
DATA FORMATS FOR OBJECTS

Express Mail Label No.

EM417217815US

Only for nonprovisional applications under 37 CFR § 1.53(b)

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

ADDRESS TO:

Box Patent Application
Assistant Commissioner for Patents
Washington, D.C. 202311. ☐ General Authorization Form & Fee Transmittal
(Submit an original and a duplicate for fee processing)2. ☒ Specification [Total Pages] **42**
(preferred arrangement set forth below)

- Descriptive Title of the Invention
- Cross References to Related Applications
- Statement Regarding Fed sponsored R & D
- Reference to Microfiche Appendix
- Background of the Invention
- Brief Summary of the Invention
- Brief Description of the Drawings (if filed)
- Detailed Description
- Claim(s)
- Abstract of the Disclosure

3. ☒ Drawing(s) (35 USC 113) [Total Sheets] **23**4. Oath or Declaration [Total Pages] **2**

- a. ☐ Newly executed (original or copy)
- b. ☒ Copy from a prior application (37 CFR 1.63(d))
(for continuation/divisional with Box 17 completed)
- i. ☐ **DELETION OF INVENTOR(S)**
Signed statement attached deleting
inventor(s) named in the prior application,
see 37 CFR 1.63(d)(2) and 1.33(b)

5. ☒ Incorporation By Reference (useable if box 4b is checked)
The entire disclosure of the prior application, from which a
copy of the oath or declaration is supplied under Box 4b,
is considered to be part of the disclosure of the
accompanying application and is hereby incorporated by
reference therein.6. ☐ Microfiche Computer Program (Appendix)7. Nucleotide and Amino Acid Sequence Submission
(if applicable, all necessary)

- a. ☐ Computer-Readable Copy
- b. ☐ Paper Copy (identical to computer copy)
- c. ☐ Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. ☐ Assignment Papers (cover sheet & document(s))
9. ☐ 37 CFR 3.73(b) Statement (when there is an assignee) ☐ Power of Attorney
10. ☐ English Translation Document (if applicable)
11. ☐ Information Disclosure Statement (IDS)/PTO-1449 ☐ Copies of IDS Citations
12. ☒ Preliminary Amendment
13. ☒ Return Receipt Postcard
14. ☐ Small Entity Statement(s) ☐ Statement filed in prior application, Status still proper and desired
15. ☐ Certified Copy of Priority Document(s)
(if foreign priority is claimed)
16. ☒ Other: Certificate of Express Mail

17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment

☒ Continuation ☐ Divisional ☐ Continuation-In-Part (CIP) of prior Application No.: **08/884,448**Prior application information: Examiner **Dr. Hugh Jones**Group / Art Unit **2763**☐ Claims the benefit of Provisional Application No. _____

CORRESPONDENCE ADDRESS

Maurice J. Pirio
Seed and Berry LLP
6300 Columbia Center
701 Fifth Avenue
Seattle, Washington 98104-7092
(206) 622-4900 phone
(206) 682-6031 fax

Respectfully submitted,

TYPED or PRINTED NAME Maurice J. PirioREGISTRATION NO. 33,273

SIGNATURE

Maurice J. Pirio

Date

February 8, 1999

u:\float\rosie\2001 Cover

EXPRESS MAIL NO. EM417217815US

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Present Application:

Applicant : Antony S. Williams
Title : A METHOD AND SYSTEM FOR REGISTERING DATA
FORMATS FOR OBJECTS
Docket No. : 663005.482C3
Date : February 8, 1999

Prior Application:

Examiner : Dr. Hugh Jones
Art Unit : 2763
Application No. : 08/884,448

Box Patent Application
Assistant Commissioner for Patents
Washington, DC 20231

PRELIMINARY AMENDMENT

Sir:

Please amend the above-identified application as follows:

In the Claims:

Please cancel claims 1-26. Please add new claims 27-70.

- 1 27. A computer-based method for providing a client with data format
2 information relating to data formats supported by a server, the method comprising:
3 under control of server code, providing the data format information;

4 storing the provided data format information in a persistent registry; and
5 under control of the client,
6 retrieving from the persistent registry the stored data format information; and
7 determining from the retrieved data format information data formats supported by
8 the server.

1 28. The method of claim 27 wherein the server code is code executed during
2 installation of the server.

1 29. The method of claim 27 wherein the server code is code executed when
2 the server is launched.

1 30. The method of claim 27 including when the server supports a data format
2 that is compatible with the client, launching the server.

1 31. The method of claim 30 wherein the client is executing in a process and
2 the server is launched in a separate process.

1 32. The method of claim 30 wherein the client is executing in a process and
2 the server is launched in the same process.

1 33. The method of claim 30 wherein the client and the server exchange data
2 using a compatible format.

1 34. The method of claim 27 wherein the client determines the data formats
2 while the server is not executing.

1 35. A computer-based method for determining the data formats that are
2 supported by a server, the method comprising:

3 retrieving from a persistent registry data format information, the data format
4 information being provided by server code for storage in the persistent registry; and
5 determining from the retrieved data format information data formats supported by
6 the server.

1 36. The method of claim 35 including when the server supports a data format
2 that is compatible with the client, launching the server.

1 37. The method of claim 36 wherein the client is executing in a process and
2 the server is launched in a separate process.

1 38. The method of claim 36 wherein the client is executing in a process and
2 the server is launched in the same process.

1 39. The method of claim 36 wherein the client and the server exchange data
2 using a compatible format.

1 40. The method of claim 35 wherein the client determines the data formats
2 while the server is not executing.

1 41. A computer-based method for supplying data format information for data
2 formats supported by a server, the method comprising:

3 under control of server code,

4 retrieving the data format information; and

5 storing the data format information in a persistent registry so that a client
6 can retrieve the data format information from the registry and determine the data formats that are
7 supported by the server.

1 42. The method of claim 41 wherein the client determines the data formats
2 while the server is not executing.

1 43. The method of claim 41 wherein the server code is code executed during
2 installation of the server.

1 44. The method of claim 41 wherein the server code is code executed when
2 the server is launched.

1 45 The method of claim 41 including when the server supports a data format
2 that is compatible with the client launching the server.

1 46. The method of claim 45 wherein the client is executing in a process and
2 the server is launched in a separate process.

1 47. The method of claim 45 wherein the client is executing in a process and
2 the server is launched in the same process.

1 48. The method of claim 45 wherein the client and the server exchange data
2 using a compatible format.

1 49. A computer-readable medium containing instructions for causing a
2 computer system to provide a client with data format information relating to data formats
3 supported by a server, by:

under control of server code, providing the data format information;
storing the provided data format information in a persistent registry; and
under control of the client,

retrieving from the persistent registry the stored data format information;

and

determining from the retrieved data format information data formats
supported by the server.

50. The computer-readable medium of claim 49 wherein the server code is
code executed during installation of the server.

51. The computer-readable medium of claim 49 wherein the server code is
code executed when the server is launched.

52. The computer-readable medium of claim 49 including when the server
supports a data format that is compatible with the client, launching the server.

53. The computer-readable medium of claim 52 wherein the client is
executing in a process and the server is launched in a separate process.

54. The computer-readable medium of claim 52 wherein the client is
executing in a process and the server is launched in the same process.

55. The computer-readable medium of claim 52 wherein the client and the
server exchange data using a compatible format.

56. The computer-readable medium of claim 49 wherein the client determines
the data formats while the server is not executing.

57. A computer-readable medium containing instructions for causing a
computer system to determine the data formats that are supported by a server, by:

retrieving from a persistent registry data format information, the data format
information being provided by server code for storage in the persistent registry; and

determining from the retrieved data format information data formats supported by
the server.

58. The computer-readable medium of claim 57 including when the server
supports a data format that is compatible with the client, launching the server.

1 59. The computer-readable medium of claim 58 wherein the client is
2 executing in a process and the server is launched in a separate process.

1 60. The computer-readable medium of claim 58 wherein the client is
2 executing in a process and the server is launched in the same process.

1 61. The computer-readable medium of claim 58 wherein the client and the
2 server exchange data using a compatible format.

1 62. The computer-readable medium of claim 57 wherein the client determines
2 the data formats while the server is not executing.

1 63. A computer-readable medium containing instructions for causing a
2 computer system to supply data format information for data formats supported by a server, by:

3 under control of server code,

4 retrieving the data format information; and

5 storing the data format information in a persistent registry so that a client
6 can retrieve the data format information from the registry and determine the data formats that are
7 supported by the server.

1 64. The computer-readable medium of claim 63 wherein the client determines
2 the data formats while the server is not executing.

1 65. The computer-readable medium of claim 63 wherein the server code is
2 code executed during installation of the server.

1 66. The computer-readable medium of claim 63 wherein the server code is
2 code executed when the server is launched.

1 67. The computer-readable medium of claim 63 including when the server
2 supports a data format that is compatible with the client launching the server.

1 68. The computer-readable medium of claim 67 wherein the client is
2 executing in a process and the server is launched in a separate process.

1 69. The computer-readable medium of claim 67 wherein the client is
2 executing in a process and the server is launched in the same process.

1 70. The computer-readable medium of claim 67 wherein the client and the
2 server exchange data using a compatible format.


REMARKS

Applicant has canceled claims 1-26. Claims 27-70 are pending. Applicant respectfully requests consideration of the application and its early allowance.

Respectfully submitted,

Antony S. Williams

SEED and BERRY LLP



Maurice J. Pirio

Registration No. 33,273

MJP:rap

6300 Columbia Center
701 Fifth Avenue
Seattle, Washington 98104-7092
(206) 622-4900
Fax: (206) 682-6031

u:\float\rosie\2001\prelim

Description

5 A METHOD AND SYSTEM FOR REGISTERING
DATA FORMATS FOR OBJECTS

Technical Field

10 This invention relates generally to a computer
method and system for registering data formats for objects
and, more specifically, to a method and system for storing
data formats supported by a server, retrieving data
formats supported by the server, and sending data to the
server and requesting the server to return data in the
15 retrieved format.

Background of the Invention

20 Current document processing computer systems
allow a user to prepare compound documents. A compound
document is a document that contains information in
various formats. For example, a compound document may
contain data in text format, chart format, numerical
format, etc. Figure 1 is an example of a compound
document. In this example, the compound document 101 is
25 generated as a report for a certain manufacturing project.
The compound document 101 contains scheduling data 102,
which is presented in chart format; budgeting data 103,
which is presented in spreadsheet format; and explanatory
data 104, which is presented in text format. In typical
30 prior systems, a user generates the scheduling data 102
using a project management computer program and the
budgeting data 103 using a spreadsheet computer program.
After this data has been generated, the user creates the
compound document 101, enters the explanatory data 104,
35 and incorporates the scheduling data 102 and budgeting
data 103 using a word processing computer program.

Figure 2 shows how the scheduling data,
budgeting data, and explanatory data can be incorporated

into the compound document. The user generates scheduling data using the project management program 201 and then stores the data in the clipboard 203. The user generates budgeting data using the spreadsheet program 204 and then
5 stores the data in the clipboard 203. The clipboard 203 is an area of storage (disk or memory) that is typically accessible by any program. The project management program 201 and the spreadsheet program 204 typically store the data into the clipboard in a presentation format. A
10 presentation format is a format in which the data is easily displayed on an output device. For example, the presentation format may be a bitmap that can be displayed with a standard bitmap block transfer operation (BitBlt). The storing of data into a clipboard is referred to as
15 "copying" to the clipboard.

After data has been copied to the clipboard 203, the user starts up the word processing program 206 to create the compound document 101. The user enters the explanatory data 104 and specifies the locations in the
20 compound document 101 to which the scheduling data and budgeting data that are in the clipboard 203 are to be copied. The copying of data from a clipboard to a document is referred to as "pasting" from the clipboard. The word processing program 206 then copies the scheduling
25 data 102 and the budgeting data 103 from the clipboard 203 into the compound document 101 at the specified locations. Data that is copied from the clipboard into a compound document is referred to as "embedded" data. The word processing program 206 treats the embedded data as simple
30 bitmaps that it displays with a BitBlt operation when rendering the compound document 101 on an output device. In some prior systems, a clipboard may only be able to store data for one copy command at a time. In such a system, the scheduling data can be copied to the clipboard
35 and then pasted into the compound document. Then, the budgeting data can be copied to the clipboard and then pasted into the compound document.

Since word processors typically process only text data, users of the word processing program can move or delete embedded data, but cannot modify embedded data, unless the data is in text format. Thus, if a user wants
 5 to modify, for example, the budgeting data 103 that is in the compound document 101, the user must start up the spreadsheet program 204, load in the budgeting data 103 from a file, make the modifications, copy the modifications to the clipboard 203, start up the word
 10 processing program 206, load in the compound document 101, and paste the modified clipboard data into the compound document 101.

Some prior systems store links to the data to be included in the compound document rather than actually
 15 embedding the data. When a word processing program pastes the data from a clipboard into a compound document, a link is stored in the compound document. The link points to the data (typically residing in a file) to be included. These prior systems typically provide links to data in a
 20 format that the word processing program recognizes or treats as presentation format. For example, when the word processing program 206 is directed by a user to paste the scheduling data and budgeting data into the compound document by linking, rather than embedding, the names of
 25 files in which the scheduling data and budgeting data reside in presentation format are inserted into the document. Several compound documents can contain links to the same data to allow one copy of the data to be shared by several compound documents.

30

Summary of the Invention

It is an object of the present invention to provide a method and system for registering data formats for a server application.

35

It is another object of the present invention to provide a method and system for registering data formats

that the server application can both receive and send data in.

It is another object of the present invention to provide a method and system for a client application to
 5 determine which data formats a server application supports without launching the server application.

It is another object of the present invention to provide a method and system for transferring data between a server and client application.

10 These and other objects, which will become apparent as the invention is more fully described below, are obtained by a method and system for transferring data between a server and client application. In a preferred
 15 embodiment, the data formats that the server application supports are stored in a persistent global registry. The client application retrieves the data formats from the persistent global registry and requests the server application to supply data in the retrieved format. The
 20 server application, upon receiving the request, supplies the data in the retrieved format.

Brief Description of the Drawings

Figure 1 is an example of a compound document.

Figure 2 is a block diagram illustrating the
 25 incorporation of the scheduling data, budgeting data, and explanatory data into the compound document.

Figure 3 is a block diagram illustrating the relationships between client and server applications in a preferred embodiment.

30 Figure 4 is a block diagram illustrating the relationship between an object handler and client and server processes.

Figure 5 is a block diagram illustrating the components that comprise the object linking and embedding
 35 facilities and the communications paths.

Figure 6A is a flow diagram of a client library message dispatching routine.

Figure 6B is a flow diagram of a typical client application callback routine.

Figure 7 is an overview flow diagram of a typical function used by a client application to handle
5 waiting for an asynchronous request to complete.

Figure 8 is a schematic diagram of an object data structure.

Figure 9 is an overview flow diagram illustrating a typical input loop for an application in an
10 event-driven windowing operating system environment.

Figure 10 is an overview flow diagram for the client library routine `Query_Release_Status?`

Figure 11 is an overview flow diagram of the procedure a client application follows to open or create a
15 compound document.

Figure 12 is a flow diagram of the function `Change_Object_Format` implemented by a typical client application.

Figure 13 is a flow diagram of the client
20 library function `Enum_Formats`.

Figure 14 is a flow diagram of the client library routine `Request_Data` and the corresponding server routine.

Figure 15 is a flow diagram of the client
25 library function `Get_Data`.

Figure 16 is a flow diagram of the server application routine `Server_Get_Data`.

Figure 17 is a block diagram of the object linking and embedding used to generate the spreadsheet
30 scheduling data and the weekly reports.

Figure 18 is a flow diagram of the function `Update_Object_Contents`.

Figure 19 is a flow diagram for the client library routine `Send_Data` and corresponding server
35 routine.

Detailed Description of the Invention

The present invention provides a method in which data that is contained within a compound document can be manipulated directly by the application program that creates the data. In a preferred embodiment, an application program that creates data can store data in various formats and receive data in various formats. The application program registers the formats by storing the formats in a persistent registry. Another application program can check the persistent registry to determine which formats are supported. The other application can send or receive data in a registered format. In a preferred embodiment, an application program that creates a compound document controls the manipulation of linked or embedded data that was generated by another application. In object-oriented parlance, this data is referred to as an object. (The reference Budd, T., "An Introduction to Object-Oriented Programming," Addison-Wesley Publishing Co., Inc., 1991, provides an introduction to object-oriented concepts and terminology.) An object that is either linked or embedded into a compound document is "contained" within the document. Also, a compound document is referred to as a "container" object and the objects contained within a compound document are referred to as "containee" objects. Referring to Figures 1 and 2, the scheduling data 102 and budgeting data 103 are containee objects and the compound document 101 is a container object. Continuing with the example of Figures 1 and 2, the user can indicate to the word processor that the user wants to edit a containee object, such as the budgeting data 103. When the user indicates that the budgeting data 103 is to be edited, the word processing program determines which application should be used to edit the budgeting data (e.g., the spreadsheet program) and launches (starts up) that application. The user can then manipulate the budgeting data using the launched

application, and changes are reflected in the compound document.

In a preferred embodiment of the present invention, applications cooperate using object linking and embedding facilities to create and manipulate compound documents. An application that creates a compound document is referred to as a client application, and applications that create and manipulate containee objects are referred to as server applications. Referring to Figure 2, the project management program 201 and the spreadsheet program 204 are server applications, and the word processing program 206 is a client application. A client application is responsible for selection of the various objects within the container object and for invoking the proper server application to manipulate the selected containee object. A server application is responsible for manipulating the contents of the containee objects.

In a preferred embodiment, applications are provided with an implementation-independent Application Programming Interface (API) that provides the object linking and embedding functionality. The API is a set of routines (functions) that are invoked by client and server applications that support compound documents. These routines manage, among other things, the setup and initialization necessary for client applications to send and receive messages and data to and from server applications. The API routines are divided into a client library and a server library. The client library provides routines which invoke the correct server application to act upon a particular containee object. The server library provides routines which process requests to manipulate containee objects.

Figure 3 illustrates the relationships between client and server applications in a preferred embodiment. In this embodiment, client applications and server applications are separate processes. The client process

301 includes client application code 303, client library 304, and container object 307. The server process 302 includes server application code 308, server library 309, and containee object 311. The client process 301
5 communicates with the server process 302 through the communications channel 306. The client library and server library are dynamically linked with the client application code and server application code, respectively, when an application process is started up. One skilled in the art
10 would appreciate other architectural configurations are possible. For example, the client library and server library functions could be implemented as processes separate from the client and server processes.

The use of the client library of the present
15 invention allows a client application program to be developed independently of any particular containee object format. Indeed, a client application, in general, does not need to know anything about the contents of a containee object. It is also preferred that the libraries
20 are linked to the applications dynamically when an application process is created. The dynamic linking allows applications to be marketed without library code and to be easily linked to new versions of the library.

The client library routines typically transfer
25 requests to manipulate a containee object through the communications channel 306 to the server process 302. One skilled in the art would appreciate that the communications channel 306 could be implemented through well-known interprocess communication mechanisms that are
30 provided by various operating systems. The server process 302 responds to requests to manipulate containee objects received through communications channel 306. The server library provides routines through which the server process receives requests to manipulate data and processes the
35 requests accordingly.

An example will help illustrate the relationship between the client process 301 and the server process 302.

Referring again to Figure 1, if a user wants to edit the budgeting data 103 of the compound document 101, then the following sequence of events occurs. First, the user starts up the word processor program, which is dynamically
 5 linked to the client library. Second, the user opens the compound document for editing. Third, the user selects the budgeting data, which is a containee object, and indicates that the selected object is to be edited. Fourth, the client application invokes a client library
 10 routine for performing an action on an object passing the routine a handle (which uniquely identifies the selected object) to the object and an indicator that the action is edit. Fifth, the client library routine determines that the spreadsheet program provides the actions for the
 15 budgeting data. Sixth, the client library starts up the spreadsheet program as a server process, if it is not already started. Seventh, the word processor application sends a message to the spreadsheet program that it should edit the budgeting data. Eighth, the server library
 20 receives the request to edit and invokes a routine in the spreadsheet program for editing the data. When editing is complete, the spreadsheet routine returns to the server library. The server library sends a message to the word processor application to indicate that editing is
 25 complete. The client library receives the message and returns from its invocation. Upon return from the invocation, the word processor application knows that the editing is complete.

Typically, the start up of a server process can
 30 be a relatively slow process. There are certain situations in which it may be unacceptable to incur this overhead. For example, if a user wants to print a compound document that includes many containee objects, it may take an unacceptably long time to start up the server
 35 process for each containee object and request each server process to print the object. To ameliorate this unacceptable performance, a server application can provide

code that can be dynamically linked during runtime into the client process to provide certain functionality in a more expeditious manner. This code is called an "object handler." Object handlers provide actions on behalf of the server application so that the client library routines can avoid starting up server processes and passing messages to the server process. In the above example, an object handler could provide a print routine that the client library routines could invoke to print a containee object.

Figure 4 shows the relationship between an object handler and the client and server processes. The object handler 402 is linked into the client process address space during runtime by the client library routines. Typically, the client library 403 invokes the object handler 402 directly, and the client application code need not be aware that a handler is providing the services, rather than a server process.

Figure 5 shows the components that comprise the object linking and embedding facilities and communications paths when a compound document includes containee objects implemented by different server applications. The client process 500 contains the client application 501, the client library 502, and the object handlers "A" and "B" 512, 513. The server processes 509, 510, 511 contain the server applications 506, 507, 508 and server libraries 503, 504, 505. The client process 500 establishes communications channels with each server process 509, 510, 511. Each server process 509, 510, 511 implements a particular type of containee object of the container object that is opened by the client process 500. The client application code 501 is dynamically linked to routines provided by the client library 502. When the client process 500 is running, it communicates synchronously or asynchronously with the server processes 509, 510, 511 using the functionality provided by the client library routines. Similarly, each server process

509, 510, 511 is dynamically linked to routines provided by the server libraries. When the server processes 509, 510, 511 are running, they process messages sent from the client process 500. The client library 502 sets up message passing structures and connections to each server 509, 510, 511. When the client application 501 requests an action on a containee object, the client library 502 sends an appropriate message to the server process that implements the type of the contained object. Alternatively, if the server application has defined an object handler for the requested action, then the client library 502 invokes the object handler to perform the requested action. Figure 5 shows that two servers 506, 507 have defined object handlers 512, 513. The client library routines access the persistent global registry 514 to determine information such as which server application to use for a particular containee object and whether an object handler is defined for a server application.

In addition to the client and server libraries, the object linking and embedding facilities of the present invention provide information to client and server applications through a persistent global "registry." This registry is a database of information such as (1) for each type of object, the server application that implements the object type, (2) the actions that the each server application provides to client applications, (3) where the executable files for each server application are located, and (4) whether each server application has an associated object handler.

Communication between client and server processes occurs either synchronously or asynchronously. Synchronous communication occurs when one process sends a message to the other process and the sending process suspends activity until the other process completely processes the message. For example, when a client process wants to create a new containee object, the client process (through the client library routines) sends a message to

the appropriate server process. The client process waits until the containee object is created before continuing execution. Asynchronous communication occurs when one process sends a message to the other process and the

5 sending process continues execution while the receiving process responds to the message. Typically, when the receiving process has completed responding to the request, it sends a message indicating such completion to the sending process. For example, when a client process wants

10 a containee object saved in a compound document file, the client process sends a message to the server process. The client process can continue to execute (e.g. responding to users requests) while the server process is saving the object. When the server process has completed saving the

15 object, it send a completion message to the client process.

In a preferred embodiment, messages are passed between client and server processes using interprocess communications mechanisms provided by the underlying

20 operating system. The client and server library routines provide an interface through which the details of the interprocess communication are shielded from the client and server applications. A client or server application requests a synchronous action by invoking a library

25 routine. When the routine returns to the requesting application, then the action requested has been completed. Similarly, a client or server application requests an action to occur asynchronously by invoking a library routine. The library routine initiates the action and

30 returns to the requesting application. The requesting application can then continue executing. When the requested action is complete, a message is received which the library routines process. In a preferred embodiment, the library routines use a "callback" routine to respond

35 to the asynchronously received completion message. A callback routine is provided by the requesting application. The address of the callback routine is made

available to the corresponding library so that when the library receives an asynchronous message, the library can invoke the callback routine to process the message. One skilled in the art would appreciate that there are
 5 numerous ways to accomplish this invocation depending upon the programming language employed.

Figure 6A is a flow diagram of a client library message dispatching routine. When a client process receives a message from a server process, the message is
 10 dispatched by the client library to the appropriate routine for processing. This dispatching occurs by invoking the `Process_Client_Lib_Message` routine. This routine determines to which object the message is directed and invokes the client callback routine to process the
 15 message when required. In steps 601 through 604, the client library receives a message, processes the message, and optionally invokes a client callback routine to process the message. In step 601, the routine determines for which containee object the message is directed. In
 20 step 602, the routine performs the client library provided processing. In step 603, if client application processing is needed, then the routine continues at step 604, else the routine returns. Client application processing is typically needed to respond to asynchronous messages. In
 25 step 604, the routine invokes the client application callback routine indicating the type of message received. The routine then returns.

Figure 6B is a flow diagram of a typical client application callback routine. This callback routine
 30 invokes the appropriate subroutine defined by the client application to process the asynchronous message. In steps 609 through 617, the callback routine determines the type of message and invokes the appropriate code to process the message. The client callback routine is passed a handle
 35 to the object (`↑object`) to which the message is directed and the message type. A client application callback routine typically processes the following messages:

- OBJ_CHANGED
- OBJ_CLOSED
- OBJ_SAVED
- 5 • OBJ_RELEASE

Server applications also provide callback routines that are invoked by the server library. Typically, all messages that a server receives are
10 asynchronous. A server callback routine typically processes the following messages:

- OBJ_OPEN
- OBJ_CLOSE
- 15 • OBJ_SAVE
- OBJ_DOACTION
- OBJ_RELEASE
- OBJ_SHOW
- OBJ_UPDATE

20

In addition to being passed an object handle and message type, the server callback routine receives a parameter to indicate the action that is to be performed when the message is OBJ_DOACTION. The additional parameter may
25 also be used to point to arbitrary data. The server library can then pass more than three parameters to the callback routine.

In a preferred embodiment, only one asynchronous action can be pending at any one time for any one object.
30 This restriction ensures that the response received from the client library and dispatched to the client application will be associated with a certain action request such that no computation is needed to decipher which action request the asynchronous message is
35 responding to. One skilled in the art would appreciate that sending additional information during message passing would eliminate the need for this restriction.

The client application requests an asynchronous action by invoking a client library routine. The client library routine, after initiating the request, returns the message, OBJ_WAIT_FOR_RELEASE. At that point, the client application waits to receive an OBJ_RELEASE message through its callback routine before requesting another asynchronous operation upon that object. There are basically two ways a client application can wait for the asynchronous operation to complete. One way is for the client application to continue to process all messages including user input and not generate any additional asynchronous requests for that object. The other way is for the client application to ignore all user input, but continue to get and dispatch all other messages in order to allow the client library to receive and process server library messages. From a user's perspective, the latter method will cause the application to appear unresponsive until the asynchronous request completes.

Figure 7 is an overview flow diagram of a typical function used by a client application to handle waiting for an asynchronous request to complete. The function returns only when the client library says that it is no longer processing an asynchronous request on the object. In step 701, the function calls the client library routine Query_Release_Status? to determine whether an asynchronous operation has completed on the object passed as a parameter. In step 702, if Query_Release_Status? returned OBJ_BUSY, then the function continues at step 705, otherwise the function continues at step 703. In step 703, if Query_Release_Status? returned OBJ_OK, then the asynchronous operation has completed and the function can return. Otherwise, the function must handle an error first and then return; error handling is performed in step 704. In step 705, the function enters a message receive-dispatch loop which effectively blocks until there is input, and when there is, the message is retrieved and dispatched to the appropriate message

handler. As described above, user messages can either be dispatched or filtered out. If all messages are read and dispatched, then user input continues to be processed. Otherwise, the function may choose to filter out user
 5 messages by reading them and dropping them instead of dispatching them. Once the message is processed, either after a dispatch returns or after filtering the message, the function continues at step 701.

Eventually, one of the messages received in step
 10 705 will be the notification from the server library that the server is done processing the requested asynchronous operation. This message will get dispatched as part of step 705 to the client library message dispatching routine show in Figure 6A with a pointer to the object passed in
 15 as a parameter. As part of step 602 in Figure 6A, the client library will clear the flag that indicated there was an asynchronous operation underway on the object. In step 604, the client library will invoke the callback routine with the OBJ_RELEASE notification as discussed in
 20 reference to Figure 6B. The client callback routine then returns to the client library message dispatch routine which then returns to the beginning of the loop in step 701 in Figure 7. At this point, when the client application calls the Query_Release_Status? routine,
 25 Query_Release_Status? will return OBJ_OK because the asynchronous operation has completed and the function will return.

Note that, during step 705, if the client application attempts to process another asynchronous
 30 request from the user (which is not filtered) on the same containee object, the client library will return OBJ_BUSY in response to the new request. The client application can choose to loop on this call until it returns OBJ_OK, or it can display an error notification to the user.

Client Library Services:

In addition to providing support for the various communication paths, the client library provides a set of linking and embedding functions which can be used by any client application to create and maintain a compound document. The following functions are supported by the client library:

- Open_CompoundDoc
- Save_CompoundDoc
- Close_CompoundDoc
- Create_Object
- Update_Object
- Close_Object
- Delete_Object
- Display_Object
- Activate_Object
- Query_Release_Status?
- Query_Release_Error

Use of these library functions requires the client application to initialize certain data structures upon invocation which are released when the compound document is closed. For example, when a client application is started, it opens the compound document and then allocates and initializes object data structures for each embedded or linked object contained in the compound document. These object data structures are passed as parameters to the library routines. One use of these data structures is to allow the library functions to identify and invoke the appropriate callback routine when an asynchronous operation is performed. Another use of the object data structure is to store the object type so that information can be retrieved for the object from a persistent global registry. The client application should release these data structures when the compound document is closed.

Figure 8 shows a schematic diagram of an object data structure. It consists of several items of information 801, 802, and 803, and optional information 804 that the client application may provide to hold object specific information. Item 801 is a pointer to the client application callback routine for the object. The client library uses this field in its message processing routine (Figure 6A) to invoke the callback routine when notification to the client application is required. Item 802 is the object type. In object-oriented parlance, this is known as the item "class id". The CLASS_ID field 802 identifies the particular server that implements the object class. A server can store class specific information in the persistent global registry indexed by this CLASS_ID. Item 803 is a handle to the permanent storage of the object. Item 804 comprises the remainder of the object data structure and is provided by the client application when it defines a wrapper data structure that contains the required information 801 through 803. One skilled in the art would appreciate that the exact definition of the wrapper structure depends upon the programming language used.

The client library routines are invoked by the client application code in the usual course of processing user input. In an event-driven windowing system, the client application calls the appropriate library routine in response to receiving a message indicating that the user has selected a particular menu item or object on the screen.

Figure 9 is an overview flow diagram which shows a typical input loop for an application in an event-driven windowing operating system environment. It is the loop that dispatches messages to the appropriate subroutines. In step 901, the application waits for a message. When it receives a message in step 902, the application decodes the message to determine what type of input event has occurred in steps 903 through 906. Typically, for each

type of input event, the application calls a different subroutine, steps 907 through 910. These subroutines in turn may call functions implemented by the object linking and embedding libraries. For example, when a Menu Event
 5 is received in step 905, the application will invoke the subroutine that handles processing menu selections in step 909. Step 909 will eventually call routine `Activate_Object` in step 910 if the object selected is an embedded or linked object and the user selects an action
 10 to perform on the object that requires communication with the server application.

In order to understand how these library functions work, it is useful to understand how the storage of compound documents differ from the storage of ordinary
 15 documents. An important distinction is that, in a compound document, data managed by different applications resides together (e.g. in the same file); whereas in an ordinary document, all of the data is managed and must be understood by the application that created the ordinary
 20 document. Specifically, in a compound document, embedded data is stored along with the native data of the compound document even though the client application cannot interpret the embedded data. Native data refers to data in a format that an application can process directly. For
 25 example, a word processor may use a text format for its native data. Embedded data is handled by the client application that maintains the compound document as a stream of bytes because it does not know how to interpret the embedded data. For example, if a compound document
 30 contains an embedded spreadsheet range, storage for the compound document will include the actual spreadsheet data in the format native to the server application used to implement the spreadsheet data.

One skilled in the art would appreciate that
 35 there are many ways for a client application to store data it does not understand and thus provide storage for embedded objects. One way is for the client application

to store the location of the stream of bytes for each object in an object table so the client application knows how to locate and retrieve the embedded object data when needed. Using this method, the embedded data could be
 5 stored within the same file as the client application native data. Alternatively, although logically embedded within the compound document, the data for each embedded object could actually be stored in a separate file by the client application.

10 The following scenario will help illustrate use of the client library linking and embedding functionality in a typical application. When the client application (e.g. a word processing program) creates a compound document that contains text and two embedded objects, such
 15 as a graph and spreadsheet data, it first creates and initializes storage for the entire compound document. At that point, the client application has a handle to the entire document and a handle to the native application data (the text). The procedure used by the client
 20 application to open a compound document is discussed in more detail below.

In a preferred embodiment of the present invention, the user then inserts the graph and spreadsheet objects by using an "Insert Object" command on the
 25 application "Edit" menu. In response to selection of the "Insert Object" menu item, the client application presents the user with a list of the kinds of objects that can be created. The client application constructs this list from the data in the persistent global registry. Once the user
 30 selects the desired object (in this case a graph or spreadsheet object), the client application invokes the Create_Object function to create the object. The Create_Object function sends a message to the server application that implements the graph object (through the
 35 server library) to create a graph object and to allow the user to edit the object. In one embodiment, the server application returns a handle to the newly-created graph

object. Similarly, when the user inserts the spreadsheet object, the client application invokes the Create_Object function to send a message to the server application that implements the spreadsheet object to create a spreadsheet
 5 object and to allow the user to edit the object. The server application returns a handle to the spreadsheet object.

Once these embedded objects have been created, the user can edit or otherwise manipulate the objects by
 10 selecting an object and selecting an action available on the client application menu. To execute the action, the client application invokes the client library function Activate_Object. Activate_Object invokes the server application that implements the selected object and sends
 15 the server application a message indicating the selected action and a handle to the selected object. The server application can then read and write the data of the embedded object in its normal fashion. In addition, the user can "open" an embedded object by either double
 20 clicking on the object or by selecting the object and then choosing the application menu item "Open." This "Open" action will result in the client application invoking Activate_Object on the selected object with a default action, determined from the global registry, which is
 25 typically "edit." If the user then selects a different object inside the compound document, or selects any of the native text data, the client application closes the previously selected object by invoking the Close_Object client library routine. This routine sends a message to
 30 the associated server to shut itself down.

The above scenario was described assuming that the user wanted to work with embedded objects. The same steps would be used to create and manipulate the graph and spreadsheet objects if they were instead linked objects.
 35 However, the user would specify that the client application should create a linked object when the user selects the "Insert Object" command.

In addition to the client library routines already discussed, there are several other functions used by the client application to communicate with the server. These functions are Update_Object, Display_Object, 5 Query_Release_Status?, and Query_Release_Error. Update_Object is used by the client application to request an updated presentation format from the server. The presentation format is what is actually displayed in the window of the client application. The presentation may 10 become out of date, for example, if a linked object has not been updated with respect to its source data.

Display_Object is used by the client application to request an object to be redisplayed. The client library passes the request to the server application or 15 the object handler. A bounding rectangle (a display context) is passed to the server or object handler indicating the area of redisplay.

Query_Release_Status? and Query_Release_Error are used by the client application to obtain information 20 about a previously completed asynchronous server operation. The Query_Release_Status? function is used to determine whether an outstanding asynchronous call has been completed on the specified object. That is, once a client application has been returned an 25 OBJ_WAIT_FOR_RELEASE value as a result of a function call to the client library, the client application then waits until it receives a OBJ_RELEASE message through its callback routine before the client application initiates additional asynchronous actions upon the same object. 30 This waiting is accomplished by looping on a call to Query_Release_Status? until it returns OBJ_OK (see Figure 7]. Once the client application has received the OBJ_RELEASE message, it can invoke the client library routine Query_Release_Error to determine the value 35 returned by the most recent asynchronous server operation on the specified object.

Figure 10 shows an overview flow diagram for the client library routine `Query_Release_Status?` This routine takes one parameter, a pointer to an object (an `OBJ_STRUCT` as shown in Figure 8) and returns `OBJ_BUSY` if the server for the object is unavailable or `OBJ_OK` if the server can process an asynchronous operation. In step 1001, the function first ensures that the parameter is a valid pointer to an object and returns `OBJ_ERROR` in step 1002 if it is not. In step 1003, if an asynchronous operation is in progress on the object, the function returns `OBJ_BUSY` in step 1004. This occurs when a `RELEASE` message for the object has not yet been sent to the client. The client library is responsible for keeping track of any asynchronous operations it invokes on behalf of an object. In step 1005, the client library also checks to ensure the server is not busy for other reasons, for example, the server has requested the server library to postpone all linking and embedding activities. Once a `RELEASE` message has been sent to the client application, `Query_Release_Status?` returns `OBJ_OK` in step 1006.

Opening a Compound Document

Figure 11 shows an overview flow diagram of the typical procedure a client application follows to open or create a compound document. In step 1101, the client application calls the client library function `Open_CompoundDoc` to open (or create) a file for the specified compound document. In step 1102, if the `Open_CompoundDoc` returns an error, then the client application reports this to the user in step 1103 and returns to the message-dispatch loop. In step 1104, the application reads in its linked and embedded object tables. In step 1105, the client application allocates and initializes object data structures for each embedded or linked object. In step 1106, if there are any automatic links contained in the compound document, then the function continues at step 1107, else it continues at

step 1108. In step 1107, the client application calls function Update_Object for each automatic link found. Next, in step 1108, it reads in the native data and displays the data in step 1109 using its standard
 5 application specific mechanism. In the process of displaying its data, if the client application encounters an embedded object, it calls function Display_Object, which invokes the appropriate server to display the object. In step 1110, the client application checks to
 10 see if there are any manual links. If so, the client application continues at step 1111, otherwise the application returns to the message-dispatch loop. In step 1111, the application lists the manual links and allows the user to selectively update them. A manual link is one
 15 that requires the user to explicitly tell the client application to perform the update. After the desired manual links are updated, the client application returns to the message-dispatch loop.

20 Registering Data Formats

The present invention allows a client and server application to exchange data in particular formats. The formats can be defined dynamically and displayed to the user of a client application without intervention by a
 25 server application. In a preferred embodiment, the client application obtains the formats available for a particular object class from a server application supplied list of the formats stored in the persistent global registry. The server application (or an object handler) can update the
 30 list of formats (register the formats) at any time. The client application uses the registering data format capabilities to perform operations like changing the display representation of an arbitrary object on user demand and to incorporate an independently implemented
 35 server application as an engine for processing the client native data.

As discussed above, the persistent global registry is a database of information to support the object linking and embedding functions. The persistent global registry is preferably stored on a long-term storage device, such as a disk. Data formats are stored in the persistent global registry indexed by CLASS_ID. A server application that implements an object is responsible for storing the data formats that it supports in the persistent global registry. The server applications store the data formats that it supplies data in and the data formats that it can receive data.

Two examples will serve to illustrate these capabilities of the present invention. In the example of Figure 1, a user is generating a report for a certain manufacturing project. Suppose the user wishes to change the representation of the scheduling data 102, which was produced using the project management program 201 (Figure 2), from a Gantt chart to a Pert chart. (A Gantt chart displays scheduling data as a sideways barchart with a bar for each task, and a Pert chart displays such data using circles to show tasks and connects them to show dependencies and critical paths.) The client application, the word processor program 206 (Figure 2), can provide this capability to the user through a "Change Formats" command on the client application "View" menu. When the user selects "Change Formats" for the selected object, the client application determines from the persistent global registry what formats the server application for the selected object supports, displays a list of these formats to the user for selection, and then, once the user has selected a format, requests the data from the server application in the selected format and displays the new representation.

Figure 12 shows a flow diagram of the function Change_Object_Format implemented by a typical client application. This function is called by the client application when the client determines that the user has

selected "Change Formats" from its "View" menu. The function determines what formats the selected object supports, gets the data from the server application in the format the user selects, and displays the data. In steps 5 1201 through 1204, the function calls the client library routine Enum_Formats in a loop until the client application obtains a list of all the available formats for the selected object. The client library function Enum_Formats is described in detail below. In step 1201, 10 the function calls Enum_Formats with 0 as an input parameter. Function Enum_Formats returns the first available format. In step 1202, if the format returned was NULL, then all formats have been obtained and the function continues at step 1205, otherwise it continues at 15 step 1203. In step 1203, the function stores the returned format in a list. In step 1204, the function calls Enum_Formats with the format that was returned from the previous call to Enum_Formats. The function then loops to step 1202 to test the result returned from Enum_Formats. 20 Once all of the formats have been retrieved and added to the list, the test in step 1202 will cause the function to continue at step 1205. In step 1205, the function tests to see if there is anything on the list or if there is only one format available and it is the same as what is 25 currently being displayed. If either case is true, the function informs the user in step 1206 that no other formats are available and returns. Otherwise, the function continues in step 1207.

In step 1207, the function displays the list of 30 formats to the user and obtains a format selection from the user. In step 1208, the function calls the client library function Request_Data to tell the server application to deliver the data in the desired format to the client library if the server application is available 35 to do so. The function Request_Data is described below in detail. In step 1209, if function Request_Data returns OBJ_OK, then the function continues at step 1216,

otherwise the function continues at step 1210. In step 1210, if function Request_Data returns OBJ_WAIT_FOR_RELEASE, then the function continues in step 1212, else the function handles the error in step 1211 and
 5 returns. In step 1212, the function calls the client application routine Process_Wait_For_Release, shown in Figure 7, in a mode that blocks all user input. Step 1212 completes when the client library has received an answer from the server application and has notified the client
 10 application. In step 1213, the function calls the client library function Query_Release_Error to determine the result from the Request_Data call. The function Query_Release_Error allows a client application to determine whether the asynchronous operation completed
 15 successfully. In step 1214, if this result is a handle to data, then the function can continue in step 1216 to get the data. Otherwise, the function handles the error in step 1215 and returns. In step 1216, the function calls the client library function Get_Data, described in detail
 20 below, to retrieve the data in the format selected by the user. Finally, in step 1217, the function calls a routine to display the retrieved data and returns.

Figure 13 is a flow diagram of the client library function Enum_Formats. In a preferred embodiment,
 25 Enum_Formats retrieves from the persistent global registry the next format from the list of formats available for the class of the selected object. This function takes an input parameter which is the format retrieved from a previous call to Enum_Formats and a pointer to the
 30 selected object data structure. If the format retrieved from the previous call indicates a 0, then this function returns the first format in the list. In step 1301, the function determines the object CLASS_ID from the selected object data structure. In step 1302, if the format input
 35 parameter is 0, then the function continues at step 1303, else it continues at step 1304. In step 1303, the function looks up the format list corresponding to the

object CLASS_ID and returns the first format in the list, or NULL if the format list is empty. In step 1304, the function looks up the format list corresponding to the object CLASS_ID, finds the format that matches the format
 5 input parameter, and returns the next format in the list, or NULL if there are no more formats in the list. In step 1305, the function checks to see if a format was retrieved and if one was, then it is returned, otherwise the function returns NULL.

10 Figure 14 is a flow diagram of the client library routine Request_Data and the corresponding server application processing required. The Request_Data function retrieves object data in a specified format from a server application. If the server application is busy,
 15 the function Request_Data returns with a WAIT_FOR_RELEASE message and the client application should not call the function Get_Data until the callback notification has been received. A client application calls Request_Data before calling Get_Data. Function Get_Data retrieves the data in
 20 the requested format from the object. Request_Data takes two input parameters: a pointer to the selected object data structure and the requested format. In step 1401, the function determines the object CLASS_ID from the selected object data structure. In step 1402, if the
 25 requested format is registered in the persistent global registry for the object class, then the function continues at step 1403, otherwise the function returns ERROR_FORMAT. In step 1403, the function checks the registry to determine whether there is an object handler defined for
 30 the object class. If there is none, the function continues at step 1406, otherwise it continues at step 1404. In step 1404, the function invokes the object handler to satisfy the data request. In step 1405, if the handler can satisfy the request (it returned OBJ_OK), then
 35 the function returns OBJ_OK. A handler is likely to be able to satisfy a Request_Data call in situations where the requested data format is a presentation format. If

On the server side, when the server library receives the REQUEST_DATA message, it invokes the callback routine of the server application in step 1410 passing it
20 a OBJ_REQUESTED_DATA notification and the handle to the object. The server application processes the request, and the callback routine returns a handle to the data in the requested format to the server library. If an error occurs, the callback routine returns an error value
25 instead. In step 1411, the server library sends the message REQUEST_DATA_DONE to the client library passing the handle to the data or the error value returned by the server application.

Then, when the client library asynchronously receives the REQUEST_DATA_DONE message in its message handling routine (see Figure 6A), in step 604, the library invokes the callback routine of the client application passing it a OBJ_RELEASED notification. At this point, step 1212 (see Figure 12) of the client application function Change_Object_Format will complete and user input to the client will no longer be blocked. As described in Figure 12, if the value returned by the server is a handle

to data, then the client application will be free to retrieve the data using function `Get_Data`. One skilled in the art would appreciate that different mechanisms can be used to actually pass the data. Typical examples are that

5 the data can be passed in shared memory or that each process is provided support from the underlying operating system to copy the data into its own address space by passing the operating system a handle to the data.

Figure 15 is a flow diagram of the client

10 library function `Get_Data`. The function `Get_Data` checks the object to determine if it has data in the requested format and returns the data to the client application. Three parameters are passed to the function `Get_Data`: a pointer to the object data structure, the requested

15 format, and an output parameter which is a handle to the data. In step 1501, the function determines the object `CLASS_ID` from the object data structure. In step 1502, if the input format is registered in the persistent global registry for the class of the object, then the function

20 continues at step 1503, otherwise the function returns `ERROR_FORMAT`. In step 1503, the function checks to see if client library has received data from the server application in the requested format. The object data structure maintains a list of available formats for the

25 data. This list is updated when a `Request_Data` call results in a new format being received. If the test in step 1503 fails, the function returns `ERROR_BLANK` to the client application. In step 1504, the function sets the output parameter to the handle of the data in the

30 requested format and returns `OBJ_OK`.

Figure 16 is a flow diagram of the server application routine `Server_Get_Data`. This function is invoked from the server application callback routine when it receives the `OBJ_REQUESTED_DATA` notification from the

35 server library (see step 1410 in Figure 14). The function `Server_Get_Data` allocates memory and fills it with the object data in the requested format and returns a handle

to the data or an error value. The function requires two parameters: a handle to the object and the requested format. In step 1601, the function allocates enough memory to hold the data in the requested format. In step 5 1602, it locks the memory in order to provide an atomic operation. In step 1603, the function fills the memory with data from the object in the requested format. In step 1604, it unlocks the memory. Finally, it returns either the handle to the data or an error value.

10 A second example will help illustrate how a client application incorporates registering data format capabilities to use a server application as an engine for filtering its data. Suppose that the user generating the manufacturing project report in Figure 1 is actually 15 working on a project which is a subpart of a much larger manufacturing project that involves multiple project teams. Suppose further that the user's manager maintains a spreadsheet of the scheduling data of the entire manufacturing project, and each project team is 20 responsible for entering the scheduling data for its subpart into the manager's spreadsheet. Each project team is also responsible for producing a weekly status report to the manager.

Figure 17 shows how object linking and embedding 25 is used to generate the spreadsheet scheduling data and the weekly reports. Each project team uses the database program 1701 to enter scheduling data, which is stored in a database 1710, to update the manager's spreadsheet 1706, and to produce weekly reports 1703 using a report form 30 1702.

In this example, the database program allows the user to generate reports from the data in the layout specified by a report form 1702 that the user has previously created. The report form 1702 contains display 35 fields that are database queries which are filled in by the database program when the report is actually generated. The report form 1702 is a compound document

and can contain arbitrary embedded or linked objects created by other applications. On the report form 1702, the user has placed a linked object 1703, which is the scheduling data for the user's project team. The data for
 5 the linked object is stored in the manager's spreadsheet 1706.

In this example, database program 1701 provides the ability to insert an arbitrary object (here a linked spreadsheet object) and to set or update its contents with
 10 native data through two commands, "Insert Object," and "Update Object Contents," found on the "Edit" menu of the database program. When the user selects the "Insert Object" command, the database program 1701 presents the user with a list of the classes of objects defined in the
 15 persistent global registry. Once the user has selected the type of object to create and specifies whether the object should be linked or embedded, the database program 1701 creates a default object of that CLASS_ID using the standard client library routine, Create_Object. In our
 20 example, a linked default spreadsheet object is created.

To initially set the data, or at any later time when the user wishes to update the data, the user selects the command "Update Object Contents". When "Update Object Contents" is invoked on a selected object, the database
 25 program presents a form for the user to fill in with database queries specifying the data to be placed in the selected object. Once the user has completed this form, the database program 1701 causes the data of the object to be changed by invoking the client library routine
 30 Send_Data. The range in the manager's spreadsheet 1706 is updated to reflect the scheduling data that was specified in the "Update Object Contents" form.

The spreadsheet program 1705 is used by the user's manager to view the scheduling data for the entire
 35 manufacturing project in the manager's spreadsheet 1706. This spreadsheet contains the ranges of data that were actually created by the component project teams as

described above. Whenever each project team chooses to update the manager's spreadsheet using Update Object Contents, the manager will have an updated synopsis of project progress.

5 Figure 18 shows the flow diagram for function Update_Object_Contents. This function allows the user to specify which data is to be sent to the object. In a preferred embodiment, this function is passed an object that is to have its data set. The function determines the
10 object class and retrieves from the persistent registry which data formats the object supports for setting data. The function determines if it can support any of these data formats (e.g., standard spreadsheet format). If it can, then the function allows the user to specify which
15 data from the database is to be sent to the object. In step 1801, the function displays a standard input selection form for a format that the object server supports (e.g., a spreadsheet). In step 1802, the function inputs the user input selections. In step 1804,
20 the function retrieves the data from the database as indicated by the user selections. In step 1805, the function puts the data in a format that is compatible with the object server. In step 1806, the function invokes the function Send_Data to send the formatted data to the
25 object. In step 1807, if function Send_Data returns an OBJ_OK message, the function returns, else the function continues as step 1808. In step 1808 through 1813, the function waits for the asynchronous invocation of function Send_Data to complete and the function returns. This is
30 the same process as is described in steps 1110 through 1115 of Figure 11.

Figure 19 shows the flow diagram for the client library routine Send_Data and the corresponding server processing required. The Send_Data function checks to
35 ensure the server for the selected object can set the object data in the requested format and sends the data to the server if it can. The function Send_Data has three

input parameters: a pointer to the selected object data structure, a handle to the data, and the requested format. In step 1901, the function determines the object CLASS_ID from the selected data structure. In step 1902, if the
5 input format is registered in the persistent global registry for the object class, then the function continues at step 1903, otherwise the function returns ERROR_FORMAT. In step 1903, the function checks in the registry to see if there is an object handler defined for the object. If
10 there is none, the function continues at step 1906, otherwise it continues at step 1904. In step 1904, the function invokes the object handler to satisfy the send request. In step 1905, if the handler can satisfy the request (it returned OBJ_OK) then the function returns
15 OBJ_OK. If the handler cannot satisfy the request, the program continues at step 1906. In step 1906, the function determines the location of the server for the object class. In step 1907, the function determines whether the server is connected (open). If the server is
20 not open, the function returns ERROR_NOT_OPEN, otherwise, it continues at step 1908. In step 1908, the function checks to see if the server is busy and, if it is, it returns OBJ_BUSY, otherwise it continues at step 1909. In step 1909, the function sends a SEND_DATA message to the
25 server asynchronously, passing it a handle to the data, the requested format, and a pointer to the object. Finally, the function returns OBJ_WAIT_FOR_RELEASE to the client application so that the client knows it must wait for an asynchronous response.

30 On the server side, when the server library receives the SEND_DATA message, it invokes the callback routine of the server application in step 1910 passing it a OBJ_SENT_DATA notification, a handle to the data, the requested format, and the handle to the object. If the
35 server application successfully processes the request, the callback routine will return an OBJ_OK value to the server library, otherwise the callback will return an error

value. In step 1911, the server library sends the message SEND_DATA_DONE to the client library with the value returned by the server application.

Then, when the client library asynchronously
5 receives the SEND_DATA_DONE message in its message handling routine (see Figure 6A), in step 604, the library invokes the callback routine of the client application passing it an OBJ_RELEASE notification. At this point, step 1810 (see Figure 18) of the client application
10 function Update_Object_Contents will complete.

Although the present invention has been described in terms of a preferred embodiment, it is not intended that the invention be limited to this embodiment. Modifications within the spirit of the invention will be
15 apparent to those skilled in the art. The scope of the present invention is defined by the claims which follow.

Claims

1. A method in a computer system of transferring data between a client application and a server application, the method comprising the steps of:

storing in a persistent registry a plurality of data formats that the server application supports;

under control of the client application,

retrieving and selecting a stored data format from the persistent registry; and

sending a request to the server application to supply data in the selected data format;

under control of the server application,

receiving the request to supply data in the selected data format; and

sending data in the selected data format to the client application; and

under control of the client application,

receiving the data sent by the server application.

2. The method of claim 1 wherein the step of retrieving a stored data format retrieves a plurality of data formats and the step of selecting the data format includes:

under control of the client application,

displaying the plurality of retrieved data formats on a display device; and

in response to user input, selecting a displayed data format.

3. The method of claim 2 wherein the steps of retrieving a plurality of stored data formats, displaying the retrieved formats, and selecting a displayed format are performed without launching the server application.

9. The method of claim 8 wherein the steps of retrieving a plurality of stored data formats, displaying the retrieved formats, and selecting a displayed format are performed without launching the server application.

10. The method of claim 7 wherein the steps of retrieving and selecting a stored data format are performed without launching the server application.

11. The method of claim 7 wherein the step of storing in a persistent registry a plurality of data formats is performed when the server application is loaded onto the computer system.

12. The method of claim 7 wherein the step of storing in a persistent registry a plurality of data formats is performed when the server application is launched.

13. The method of claim 7 including the steps of:
 under control of the server application,
 processing the data received from the client application; and
 sending the processed data back to the client application; and
 under control of the client application,
 receiving the processed data sent by the server application.

14. A method in a computer system for a server application to communicate to a client application the data formats it supports, the method comprising the steps of:

 storing in a persistent registry a data format that the server application supports for receiving data from the client application;

 storing in a persistent registry a data format that the server application supports for sending data to the client application; and

2025 RELEASE UNDER E.O. 14176

15. A method in a computer system of transferring data between a client application and a server application, the method comprising the steps of:

```

    storing in a persistent registry a data format that
the server application supports for sending data to the client
application;

```

under control of the client application,
retrieving and selecting from the
persistent registry a stored data format for receiving data
from the client application; and

under control of the server application,
receiving the data sent by the client
application; and

under control of the client application,
retrieving and selecting from the
persistent registry a stored data format for sending data to
the client application; and

under control of the server application,
receiving the request to supply data in
the retrieved format; and

under control of the client application,

receiving the data sent by the server application.

16. The method of claim 15 wherein the step of retrieving a stored data format for receiving data from the client application retrieves a plurality of data formats and the step of selecting the data format includes:

under control of the client application,

displaying the plurality of retrieved data formats for receiving data from the client application on a display device; and

in response to user input, selecting a displayed data format.

17. The method of claim 16 wherein the step of retrieving a stored data format for sending data to the client application retrieves a plurality of data formats and the step of selecting the data format includes:

under control of the client application,

displaying the plurality of retrieved data formats for sending data to the client application on a display device; and

in response to user input, selecting a displayed data format.

18. The method of claim 15 wherein the step of retrieving a stored data format for sending data to the client application retrieves a plurality of data formats and the step of selecting the data format includes:

under control of the client application,

displaying the plurality of retrieved data formats for sending data to the client application on a display device; and

in response to user input, selecting a displayed data format.

19. The method of claim 15 wherein the steps of retrieving and selecting a stored data format are performed without launching the server application.

20. The method of claim 15 wherein the steps of storing in a persistent registry a plurality of data formats for receiving data and storing in a persistent registry a plurality of data formats for sending data are performed when the server application is loaded onto the computer system.

21. The method of claim 15 wherein the steps of storing in a persistent registry a plurality of data formats for receiving data and storing in a persistent registry a plurality of data formats for sending data are performed when the server application is launched.

22. The method of claim 15 wherein when data is to be transferred from the client application to the server application, comprising the steps of:

- under control of the server application,
 - processing the data received from the client application; and
 - sending the processed data back to the client application; and
- under control of the client application,
 - receiving the processed data sent by the server application.

23. A method in a computer system of improving the efficiency of communication between a client application and a server application, the client application executing as a client process and having an address space, the method comprising the steps of:

- storing in a persistent registry the location of an object handler for the server application; and
- under control of the client application,

retrieving the location of the object handler for the server application;
linking the object handler into the address space of the client process; and
invoking the object handler to process a request.

24. The method of claim 23 wherein the steps of retrieving the location of and invoking the object handler are performed without launching the server application.

25. The method of claim 24 wherein the step of storing in a persistent registry the location of an object handler is performed when the server application is loaded onto the computer system.

26. The method of claim 24 wherein the step of storing in a persistent registry the location of an object handler is performed when the server application is launched.

VAC1.DOC

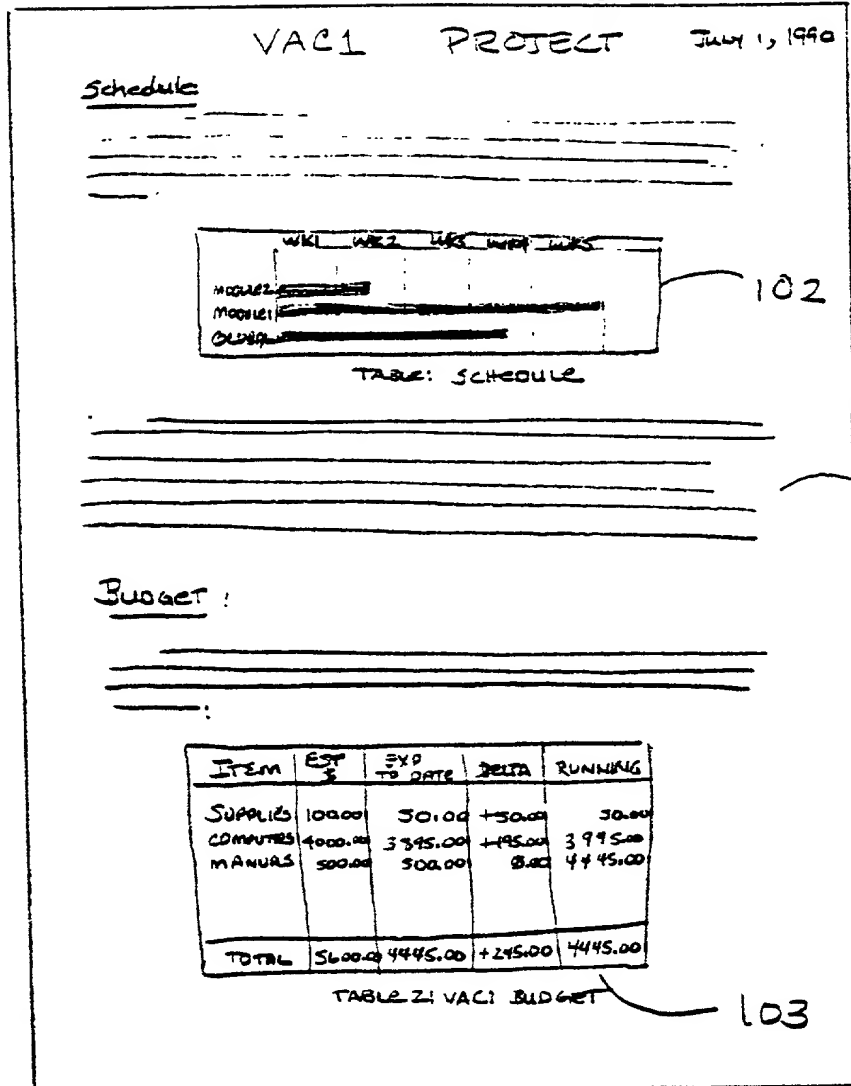


Figure 1

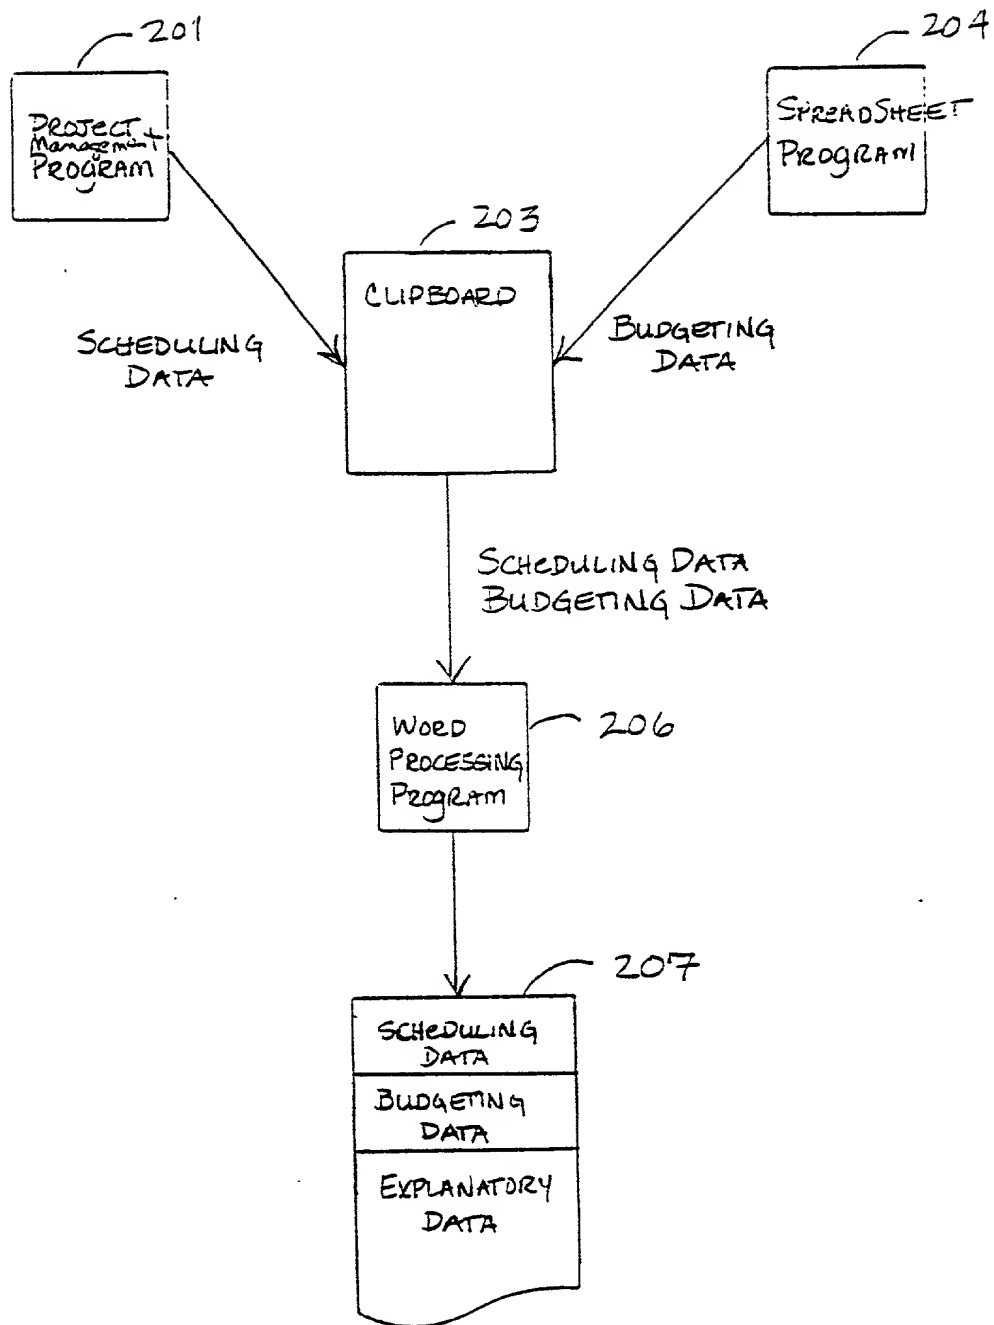


FIGURE 2




FIGURE 3

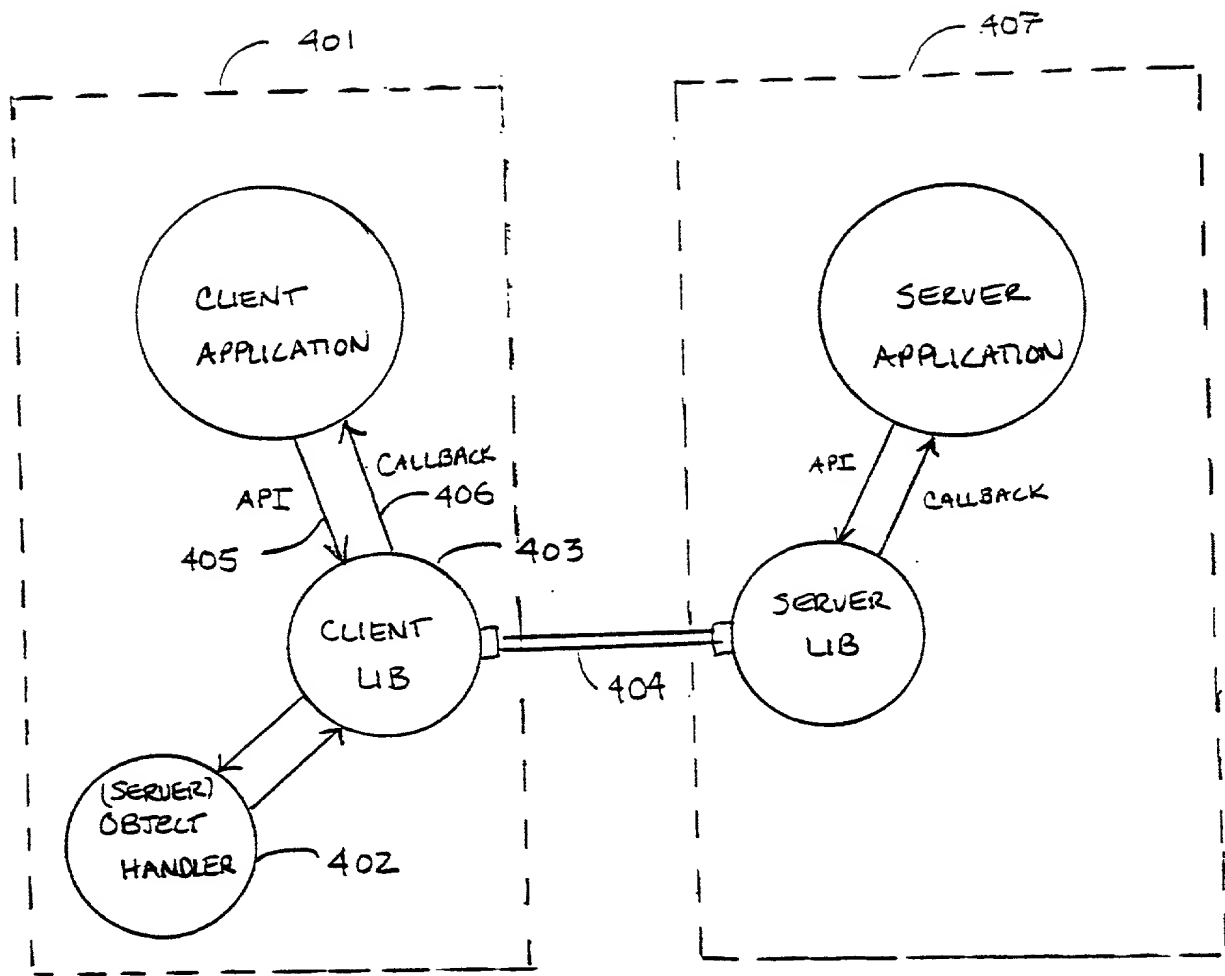


FIGURE 4

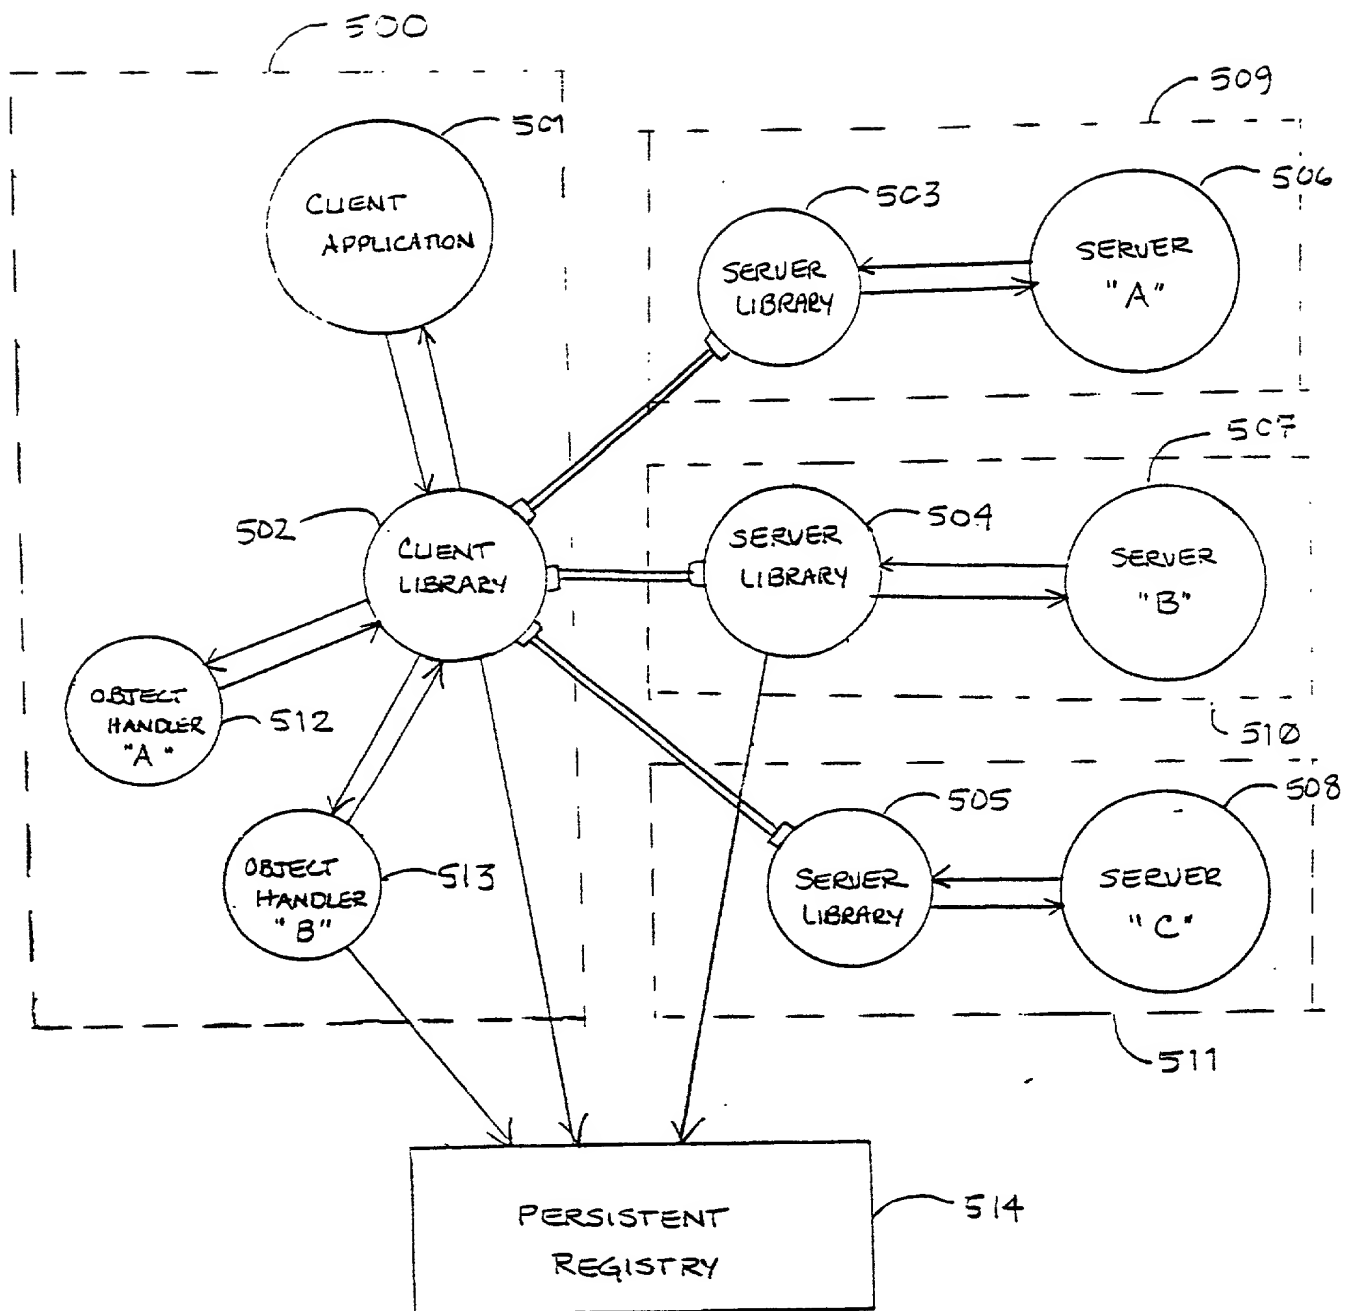


FIGURE 5




FIGURE 6A




FIGURE 6B

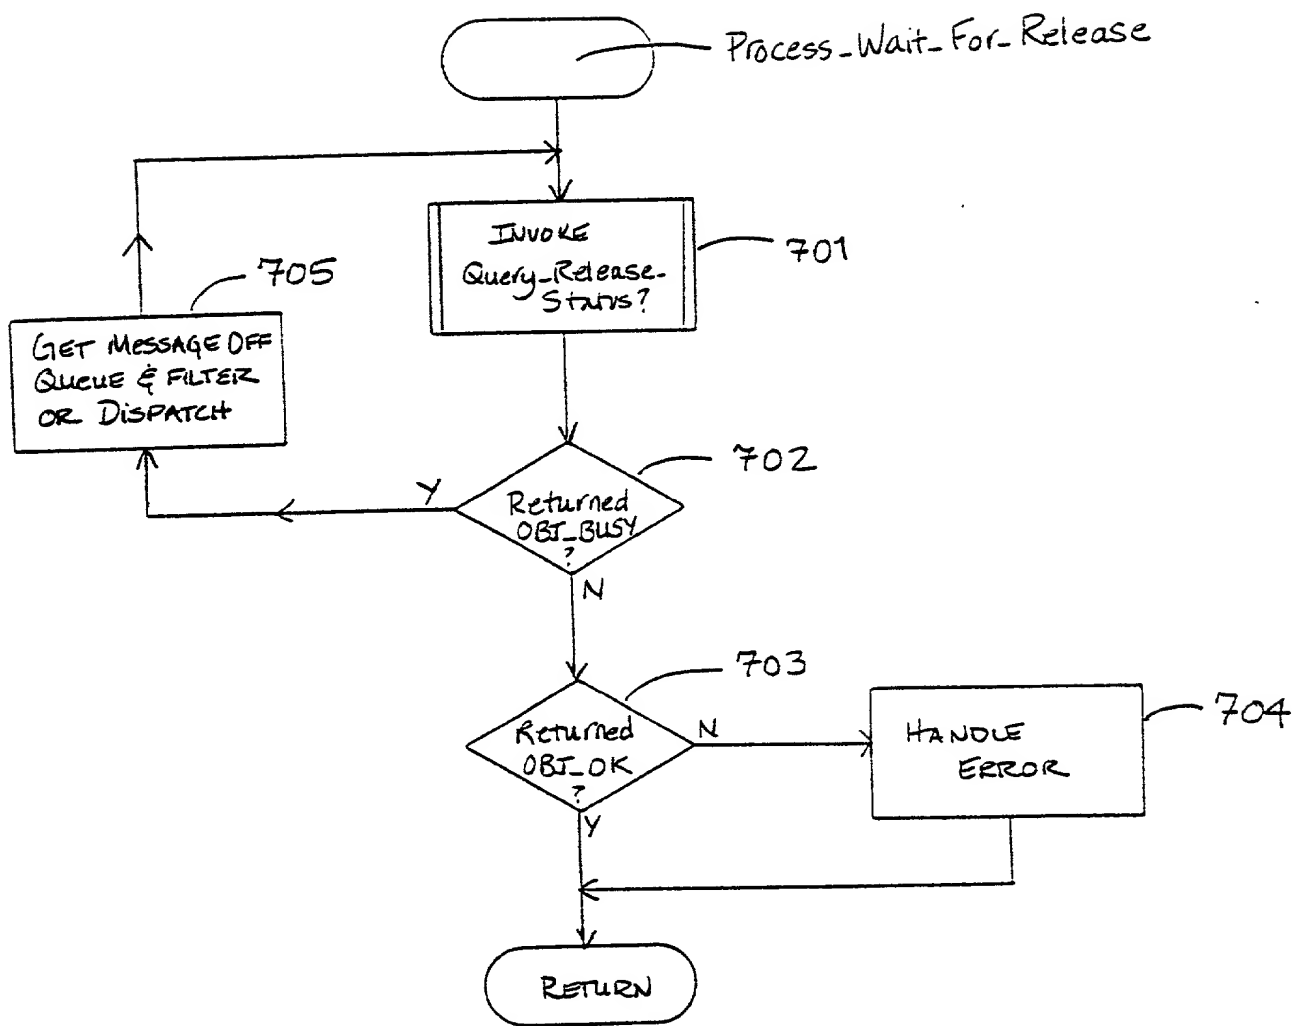


FIGURE 7

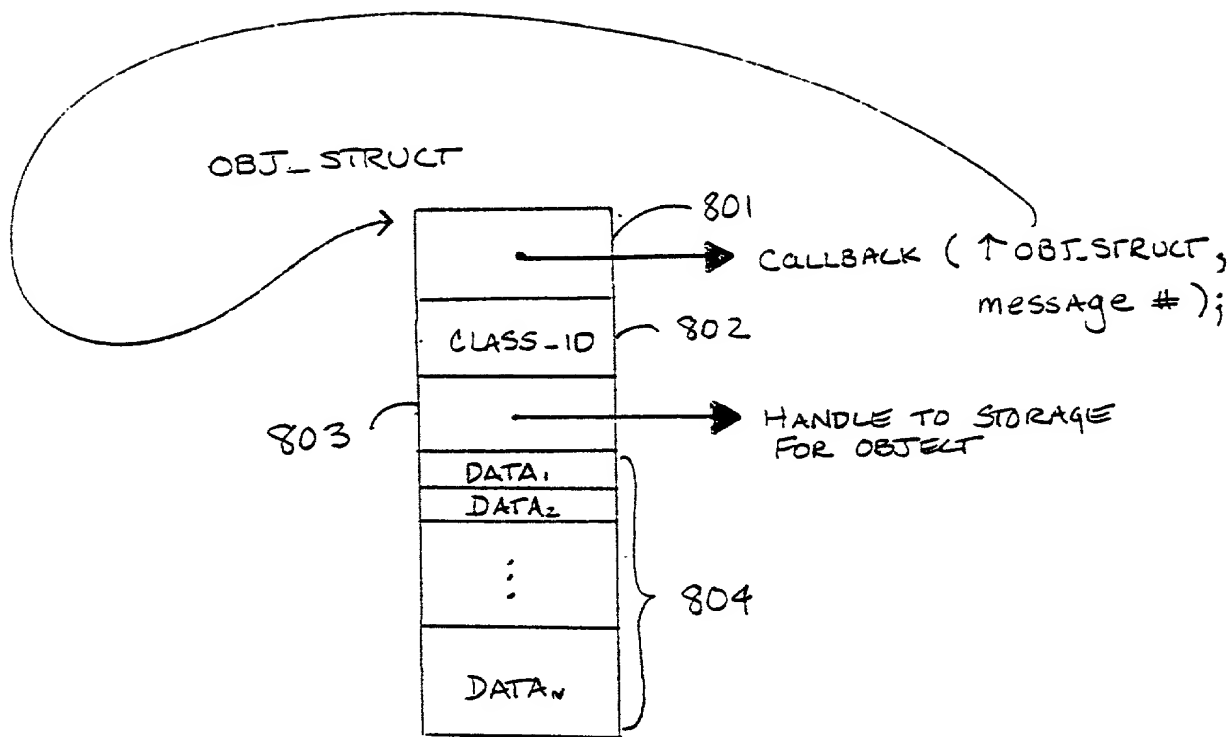


FIGURE 8

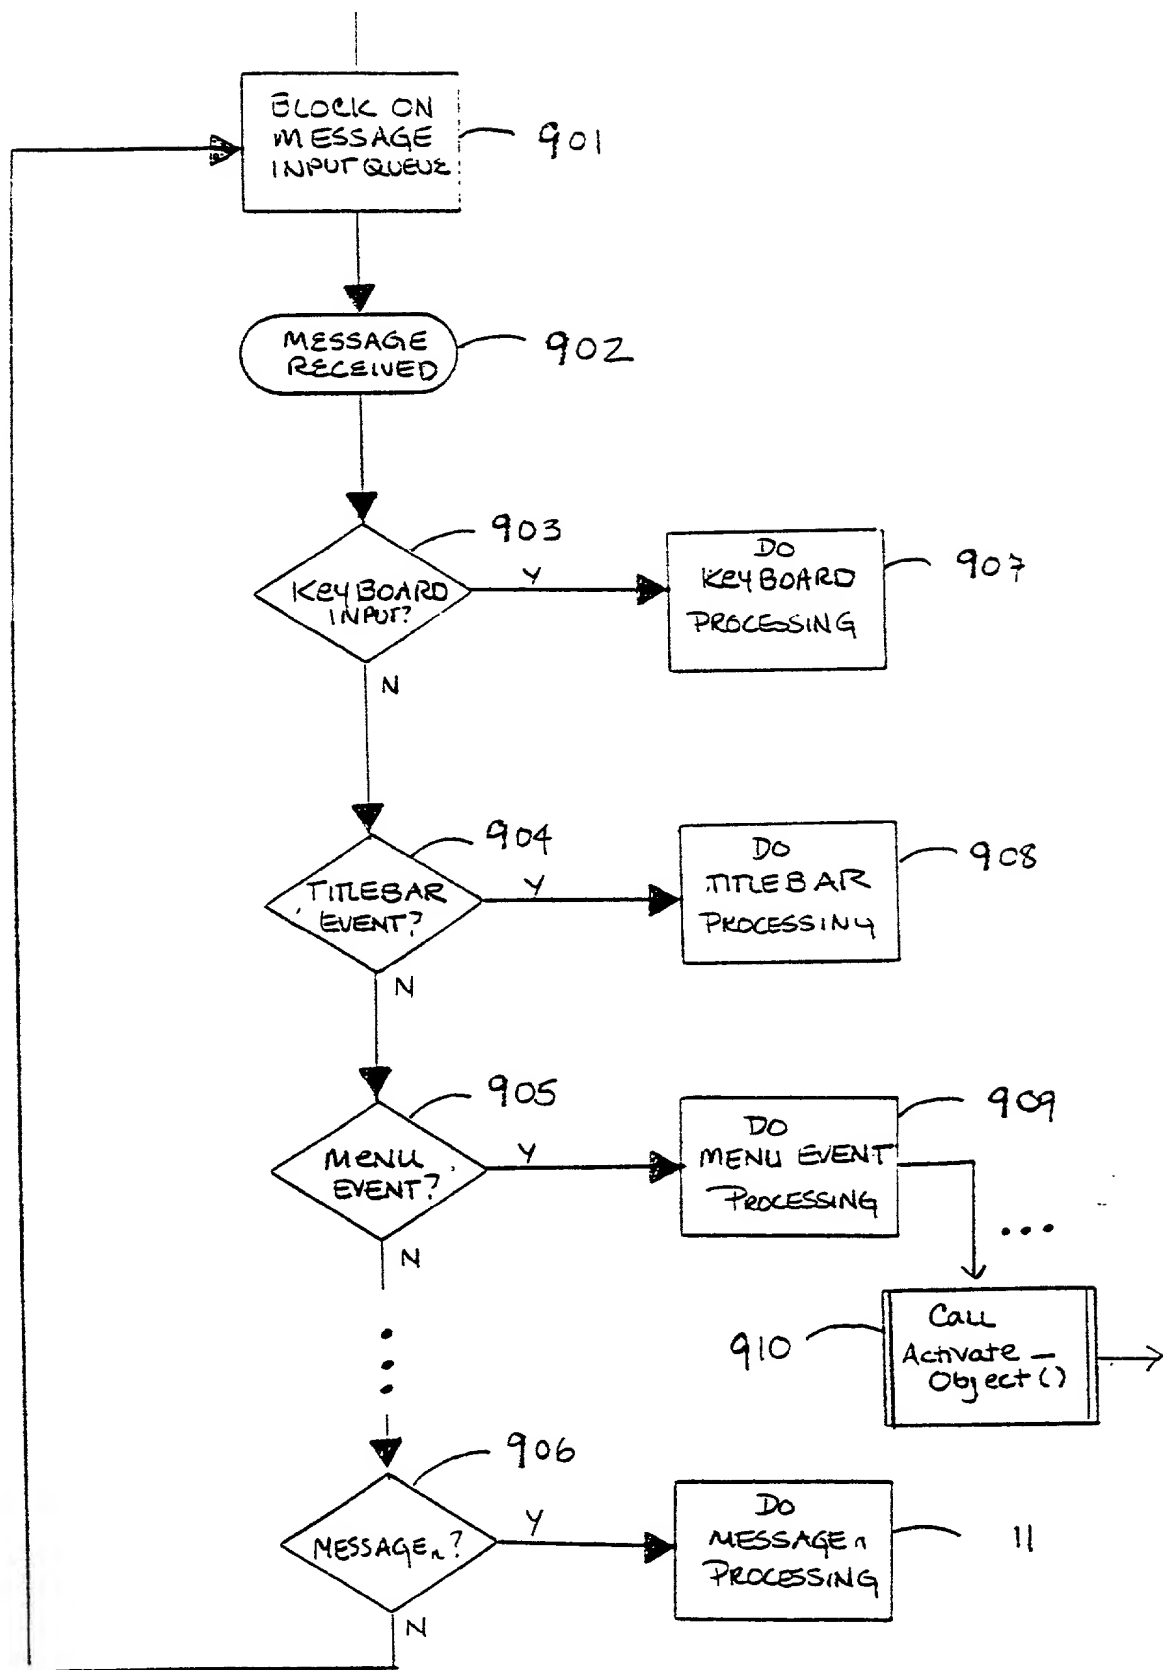


FIGURE 9

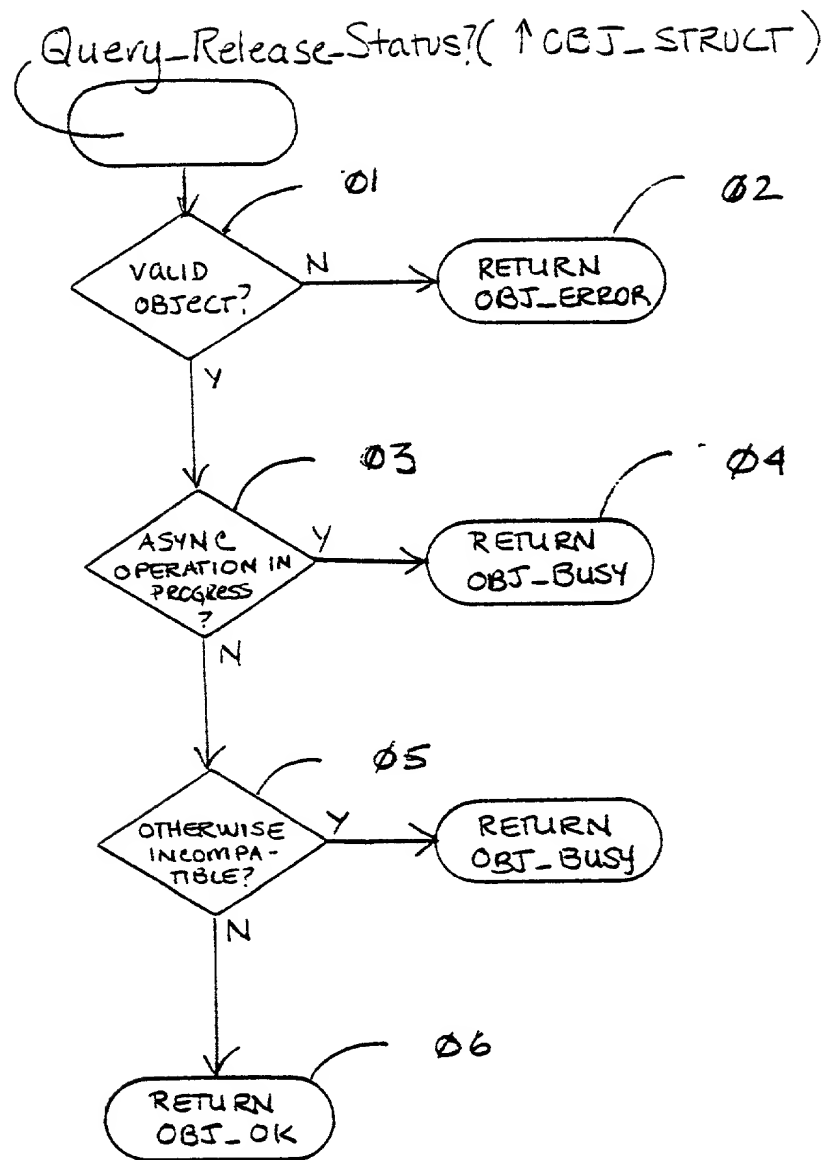


FIGURE 10

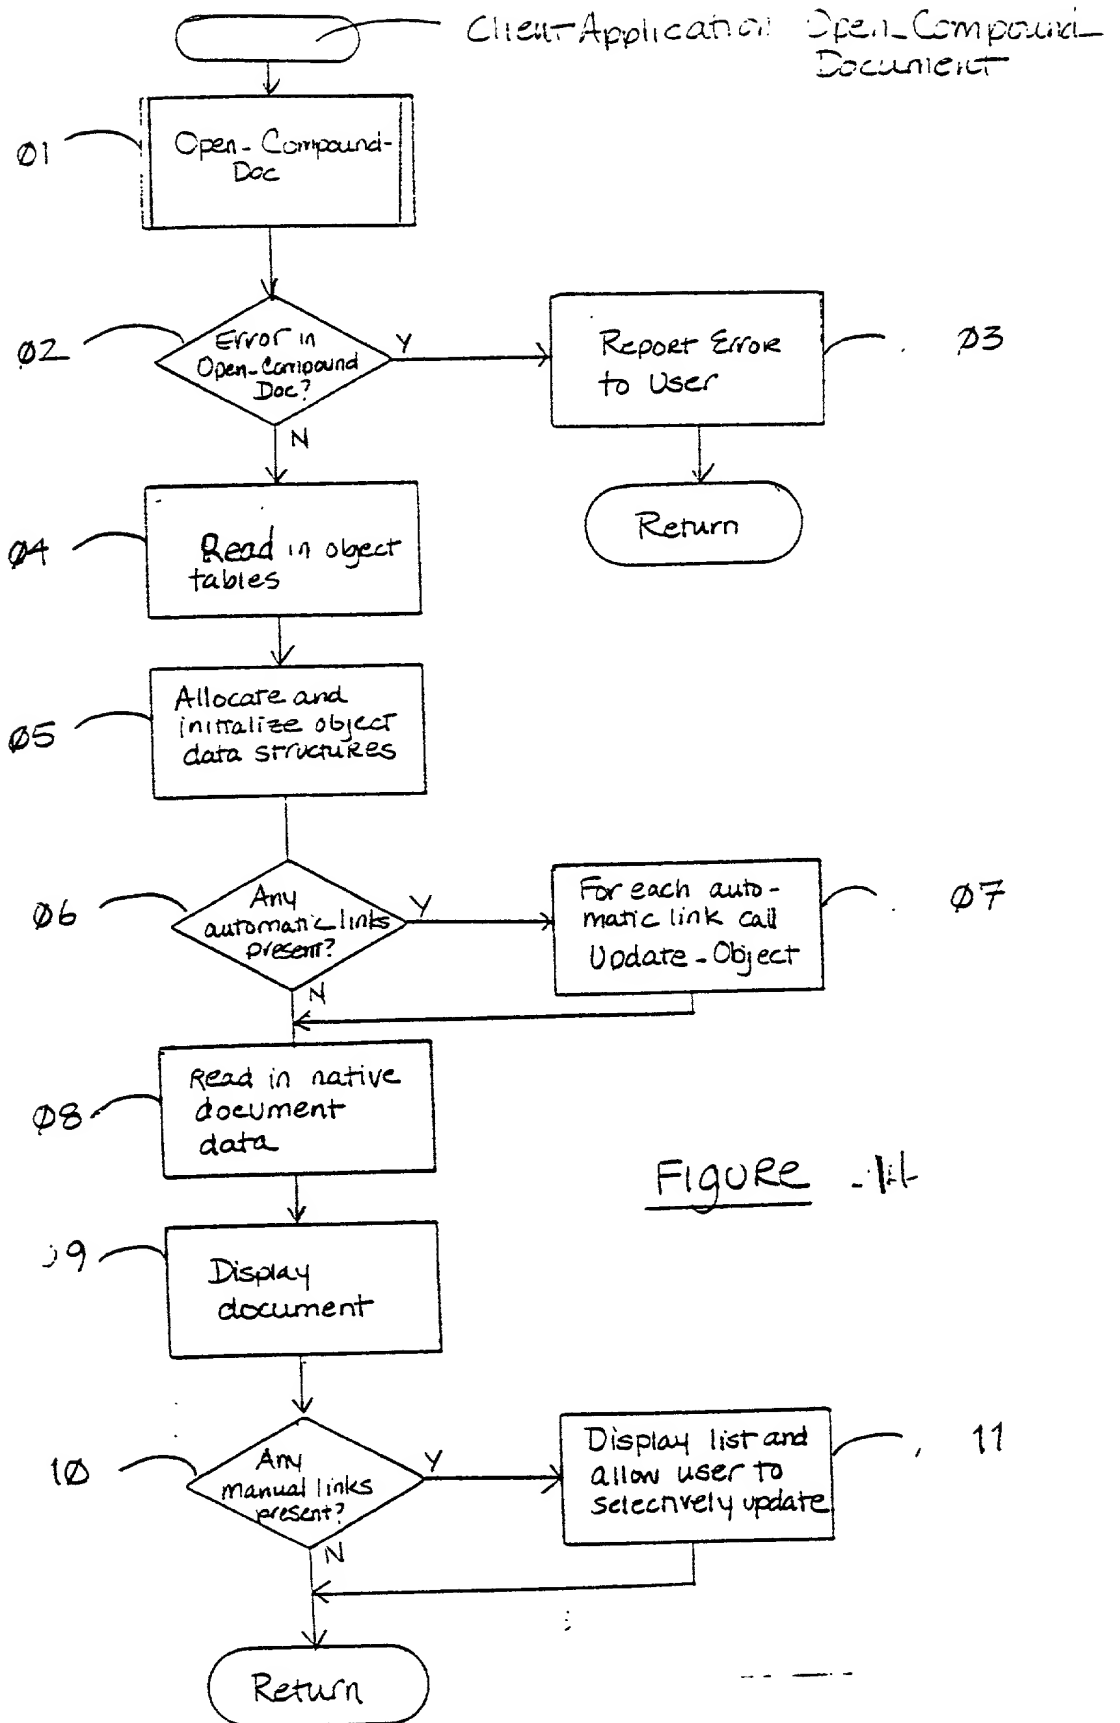


Figure -11

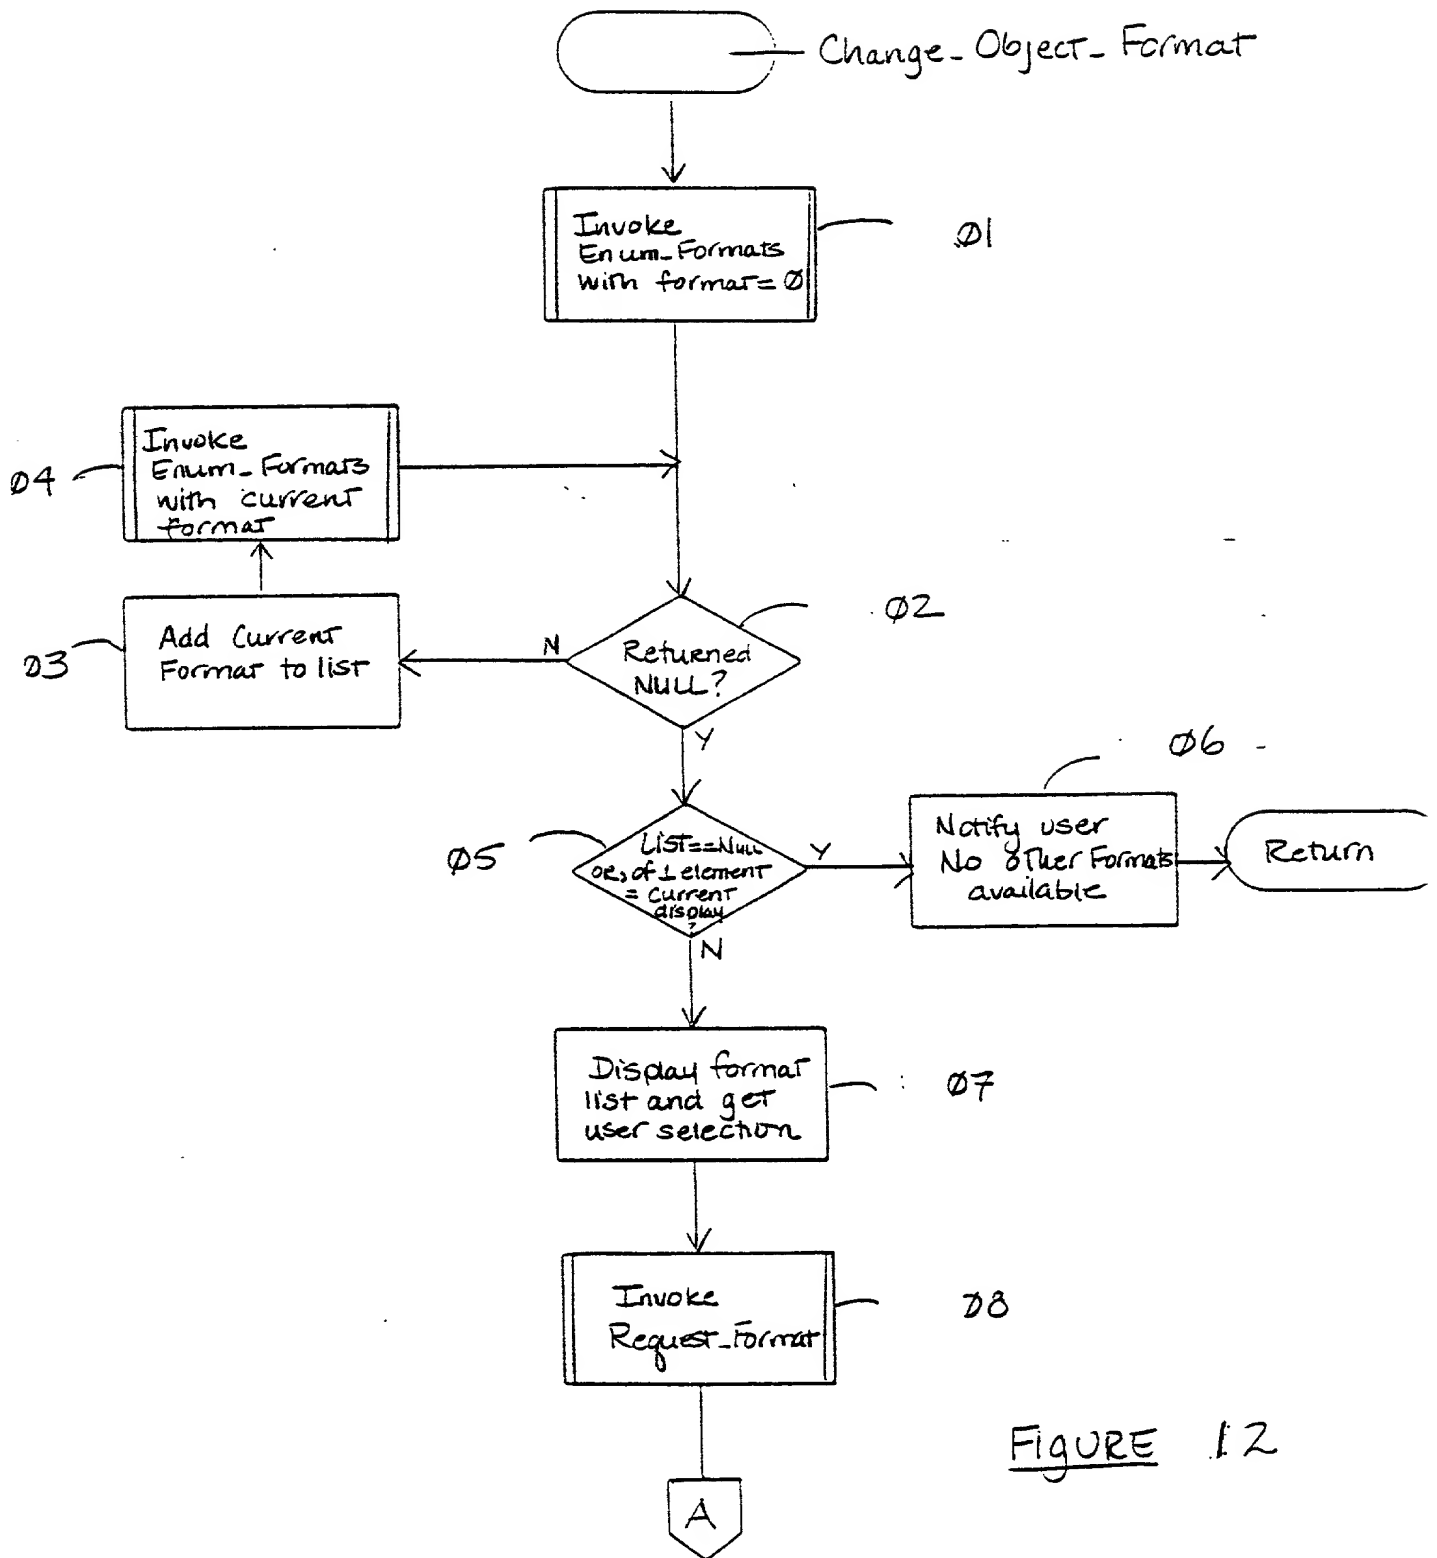


FIGURE 12

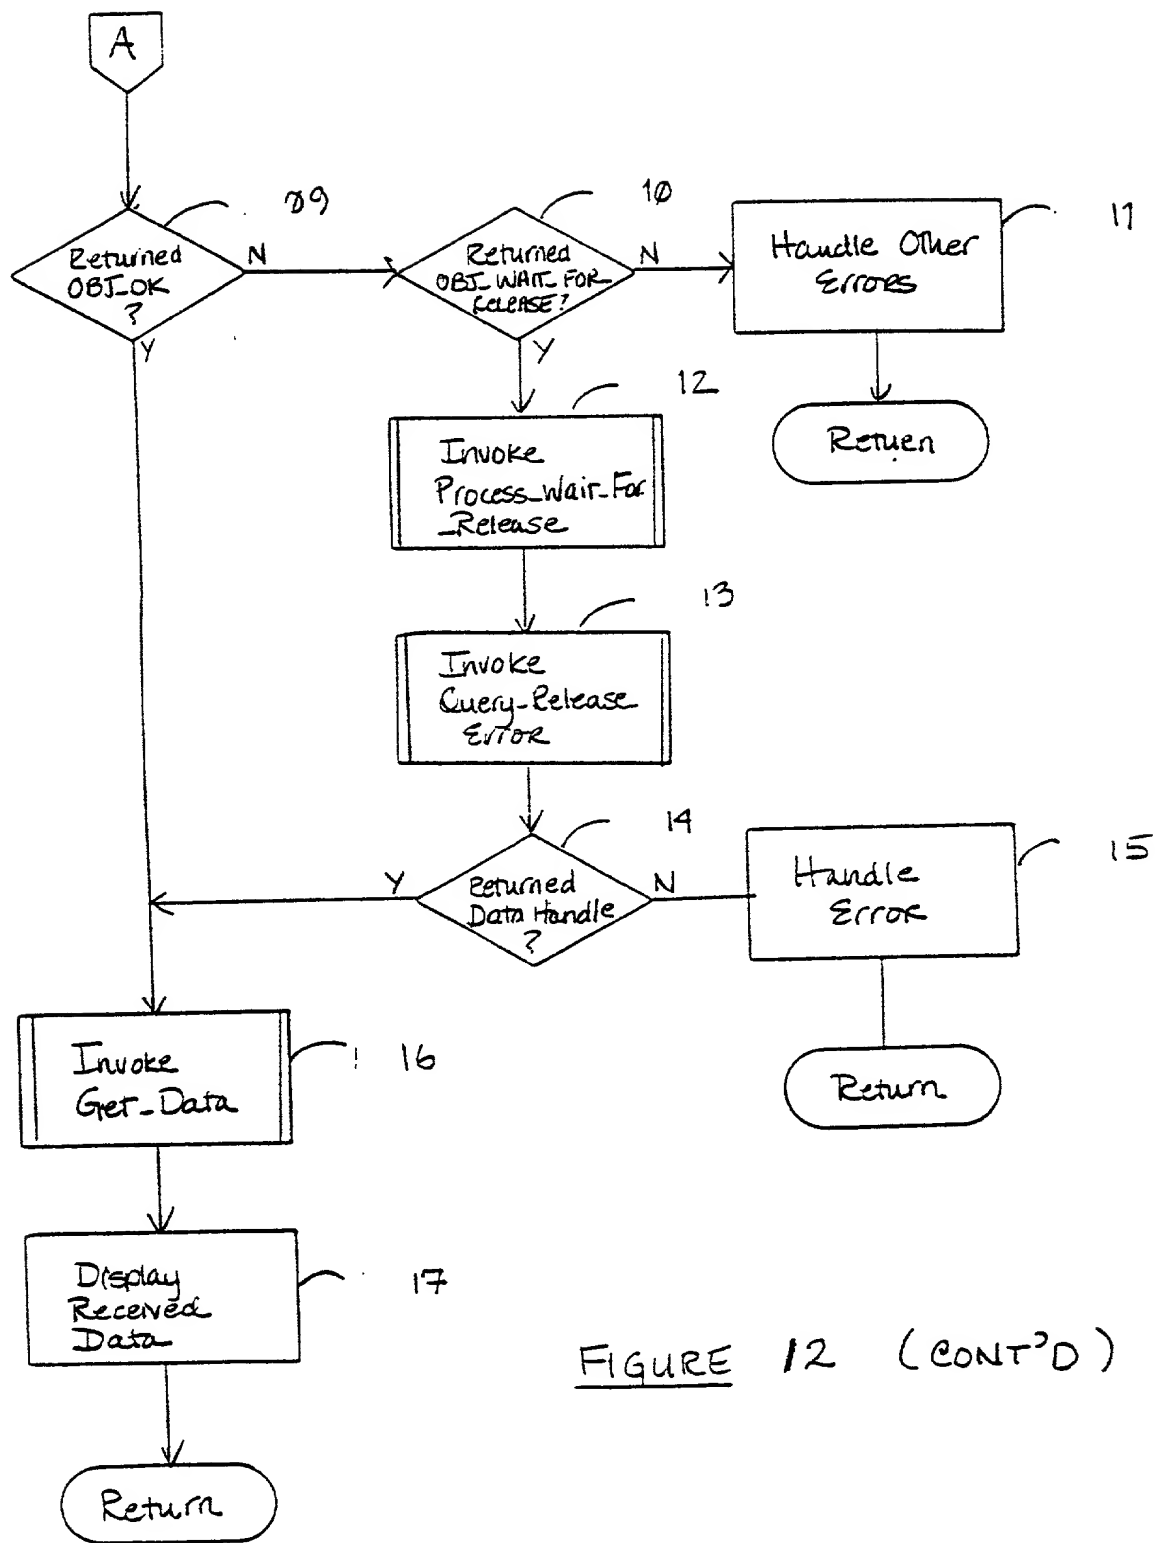


FIGURE 12 (CONT'D)

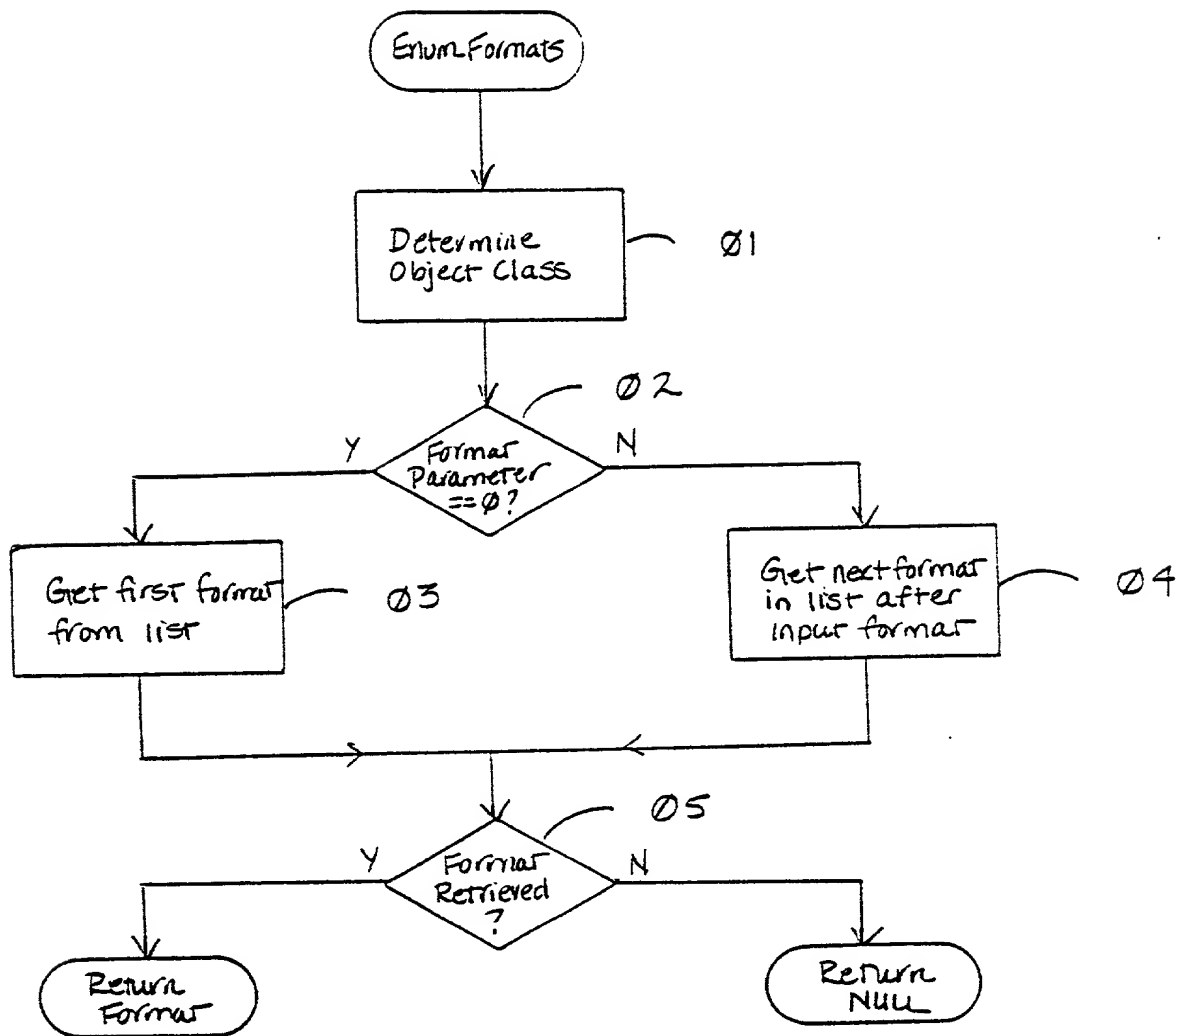


FIGURE 13

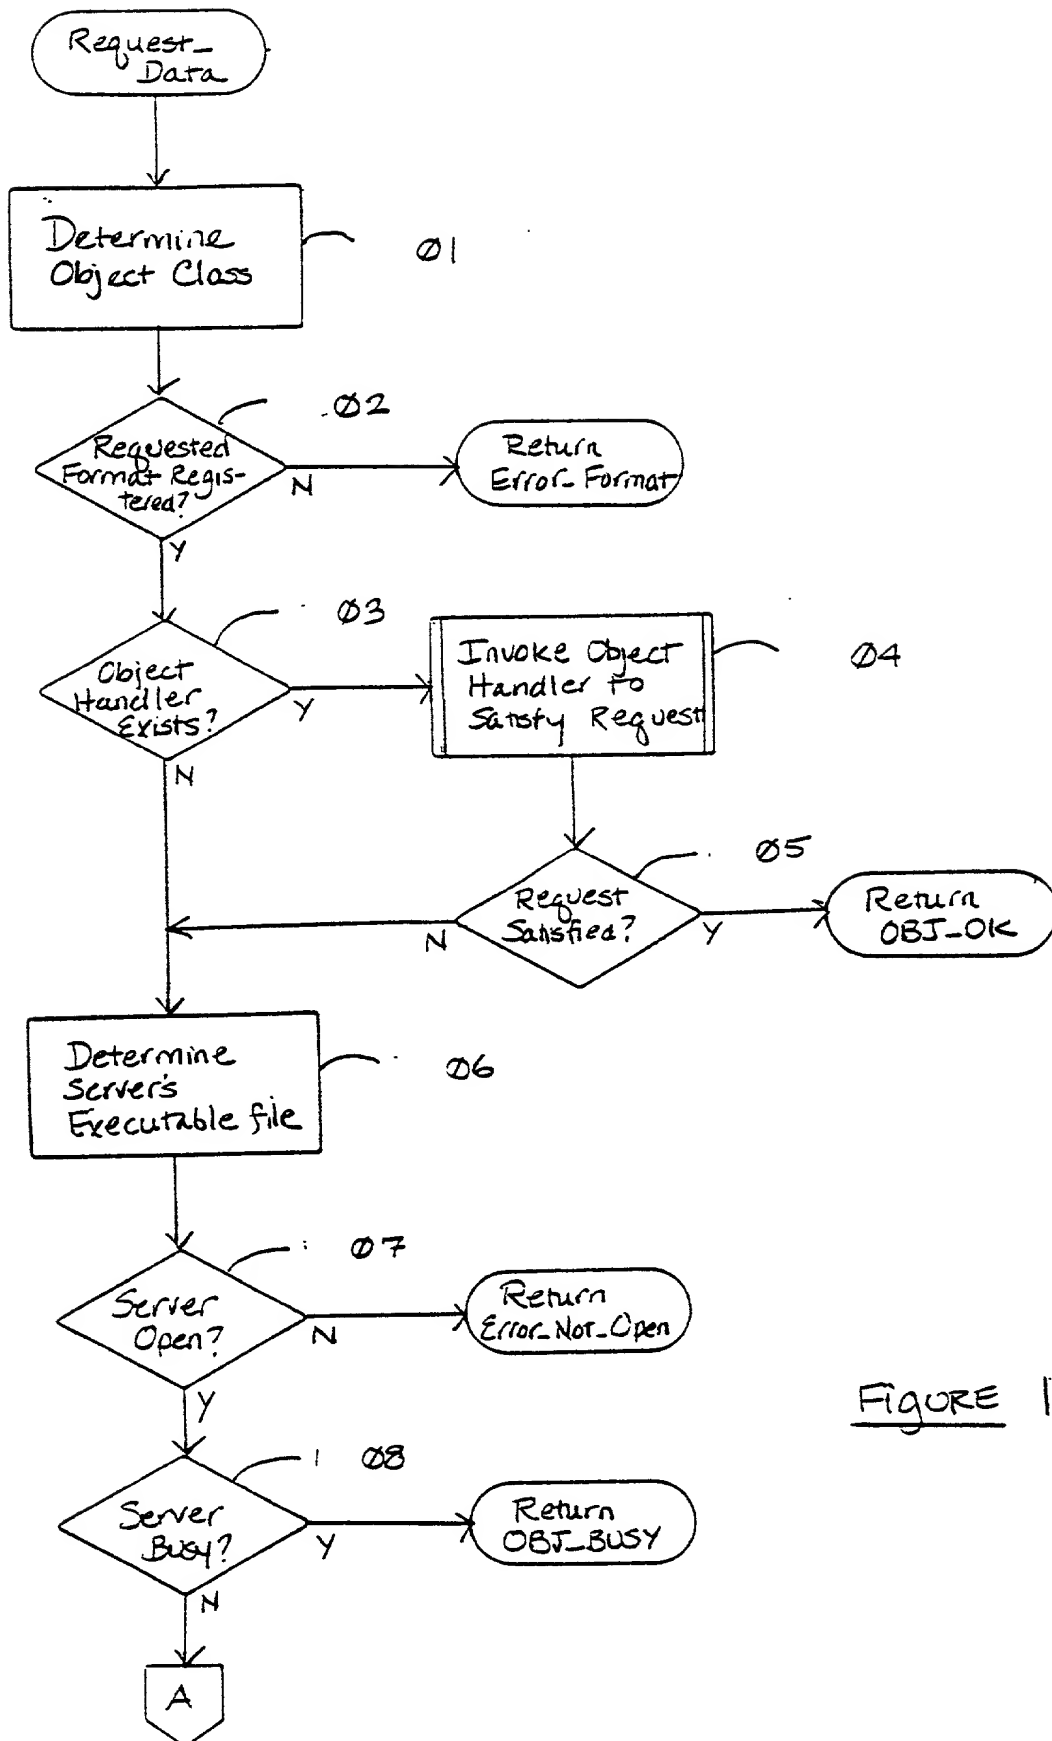


FIGURE 14

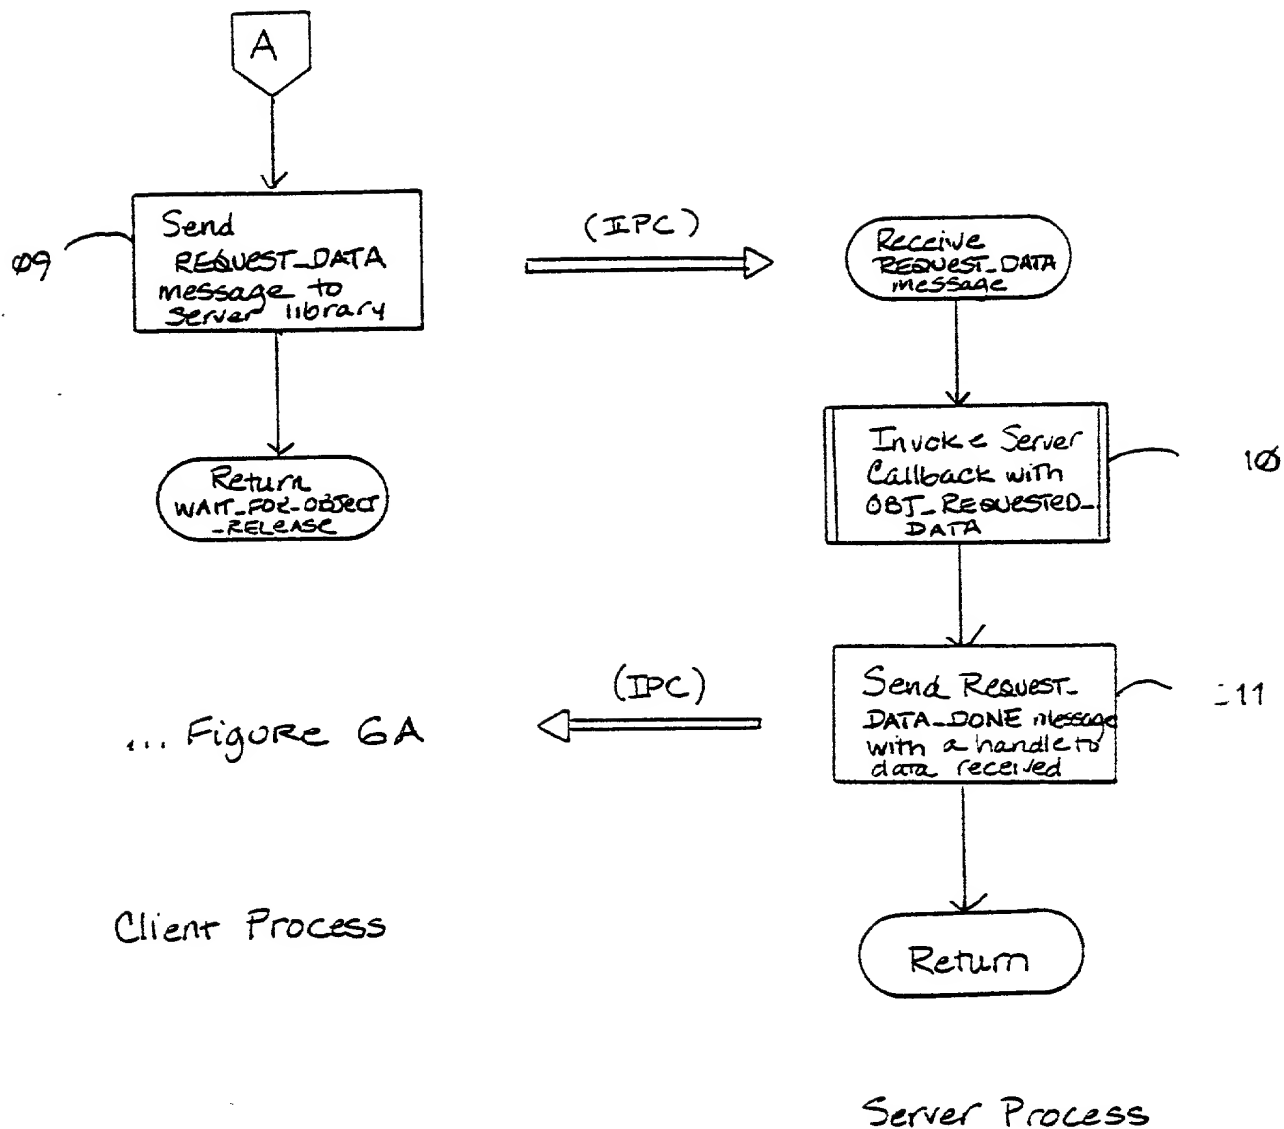


Figure 14 (CONT'D)

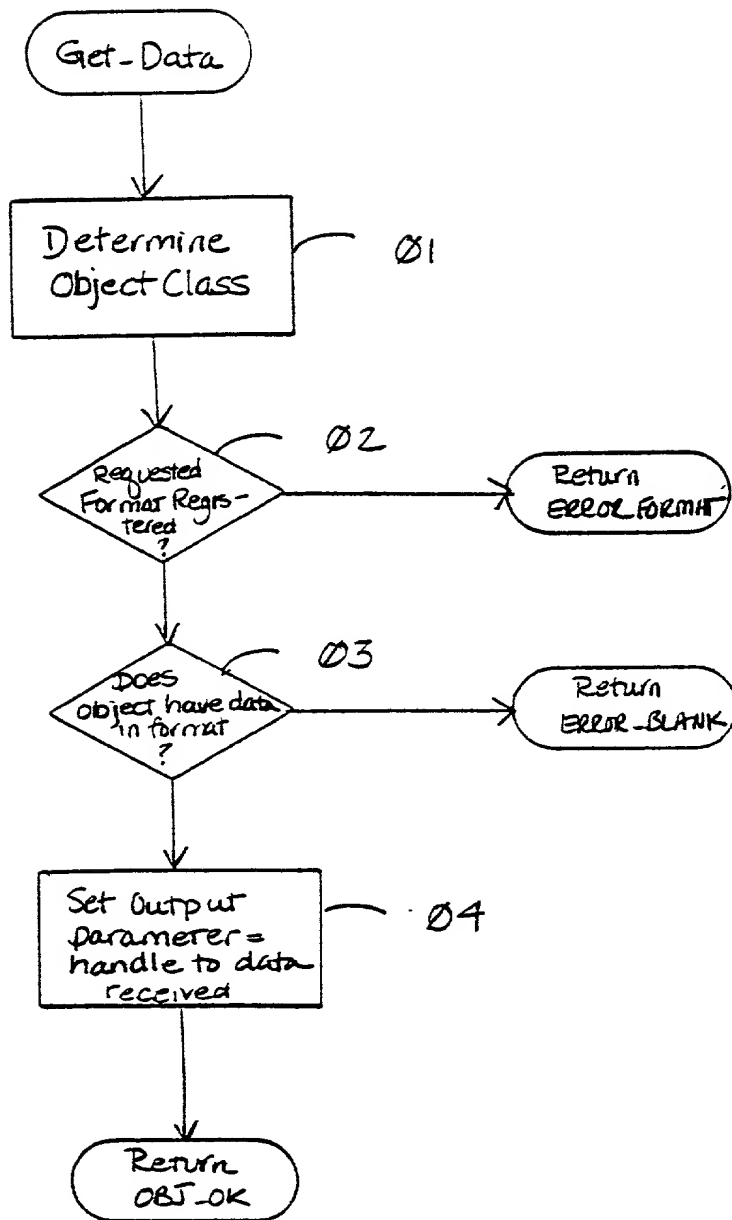


Figure 15

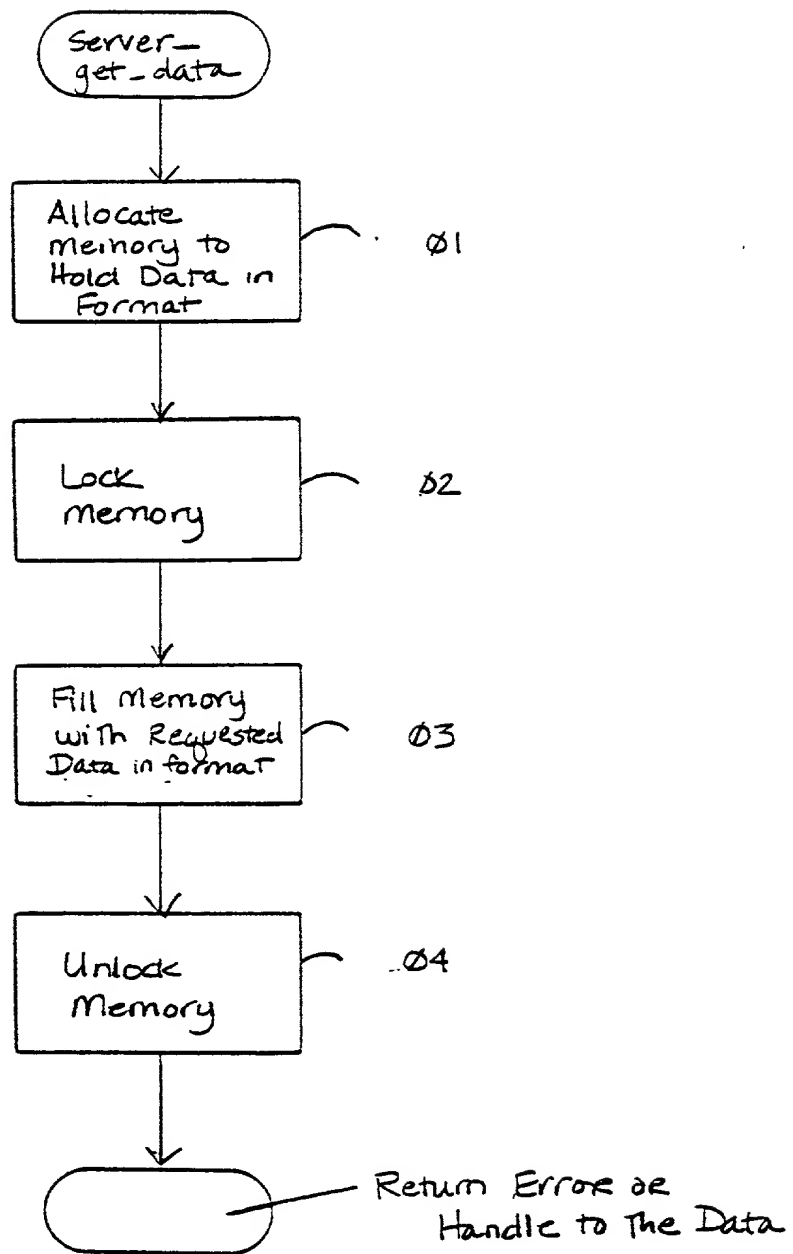


Figure 16




Figure 17.

Update-Object-Contents

Display Form for
Defining Data
& Determine
Formats

Get User
Selections

Fill in with Data
from Database
Queries

Format
Data

Send-Data

Returned
OBJ_OK?

Return

FIGURE 18

Returned
OBJ-WAIT-FOR-
RELEASE?

Invoke Process-
Wait-For-Release

Invoke
Query-Retase-
Error

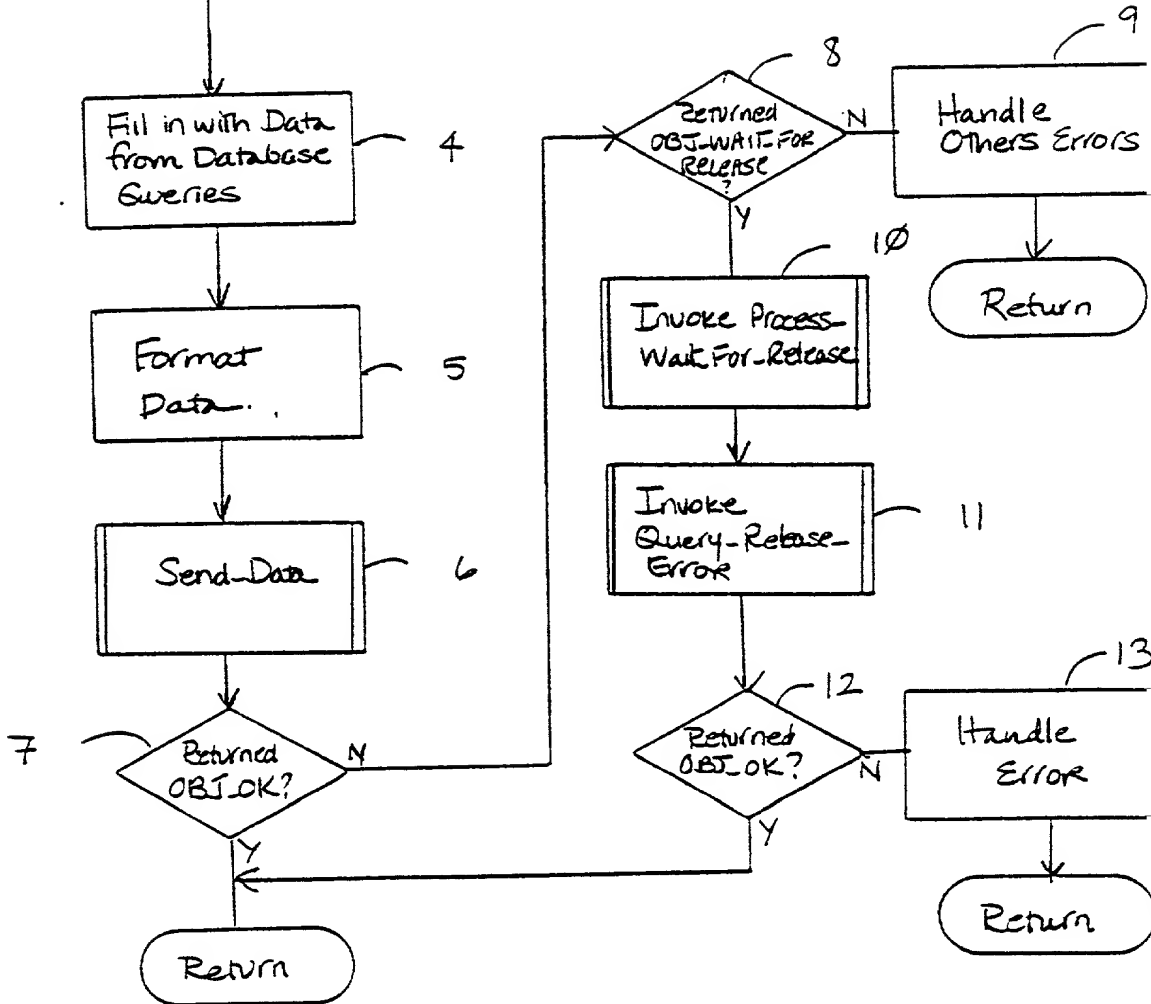
Returned
OBJ_OK?

Handle
Others Errors

Return

Handle
Error

Return



09217135-020899

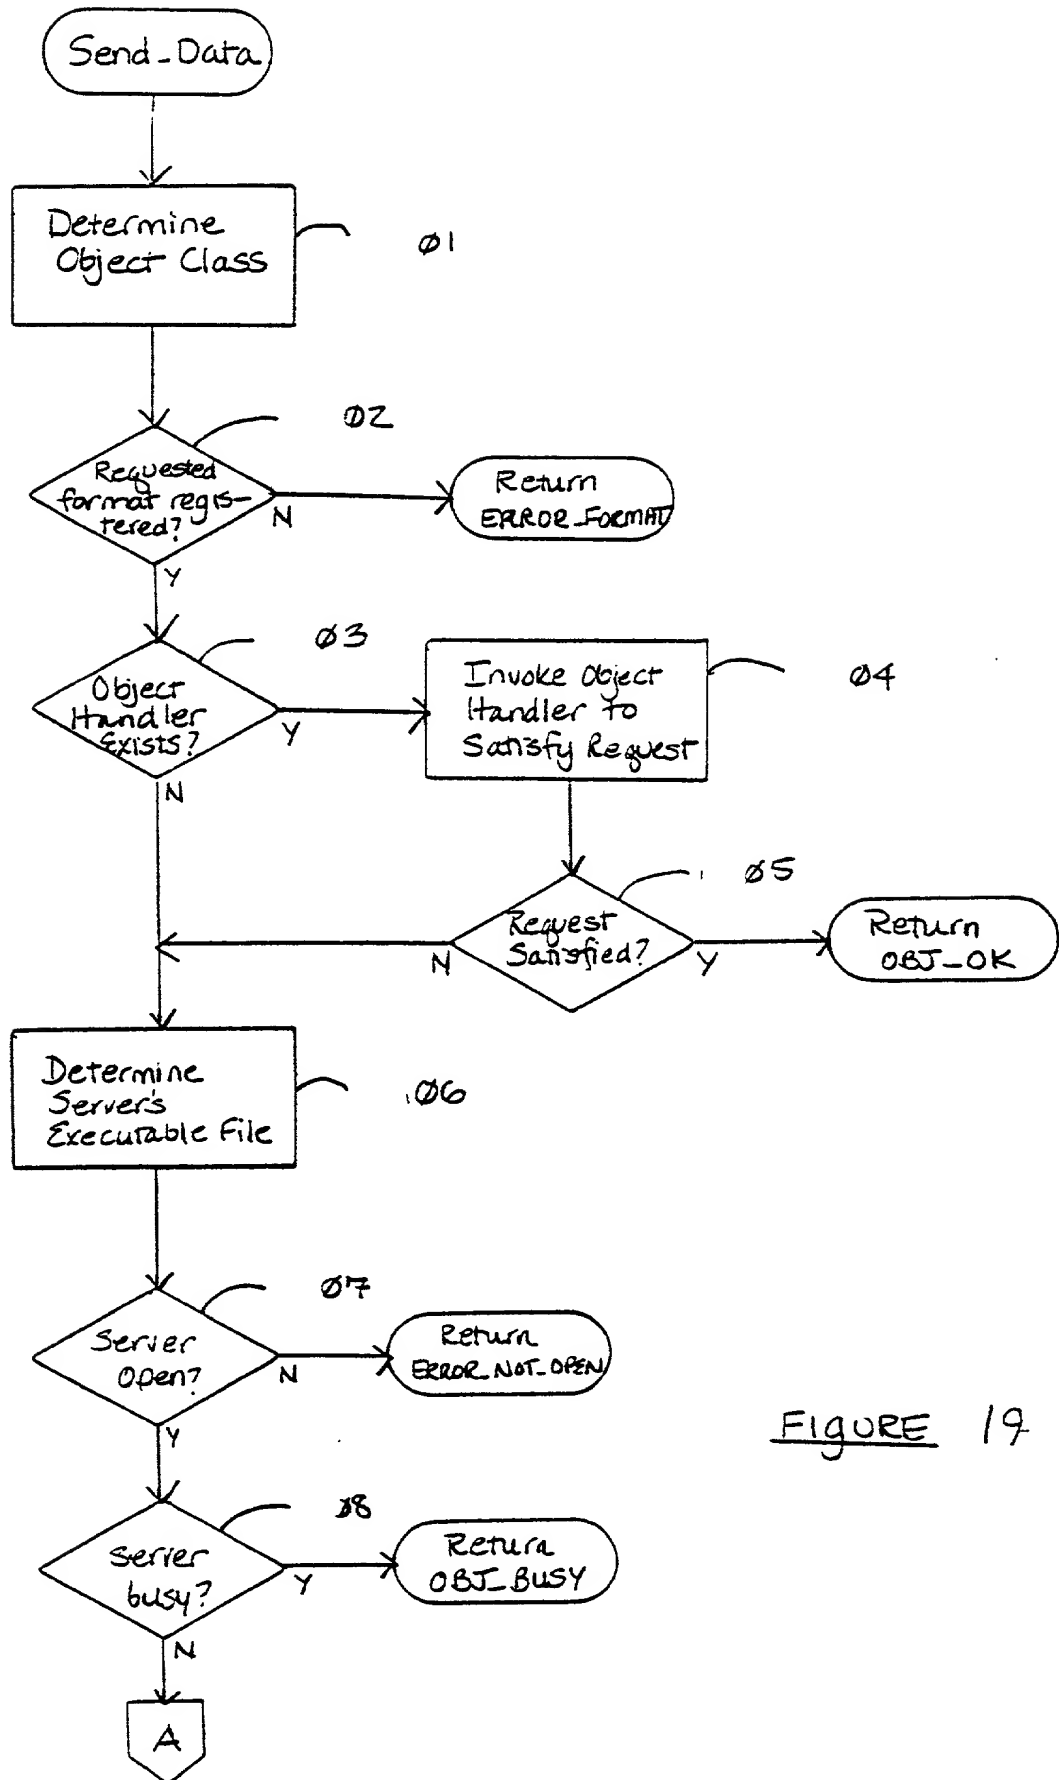


FIGURE 19

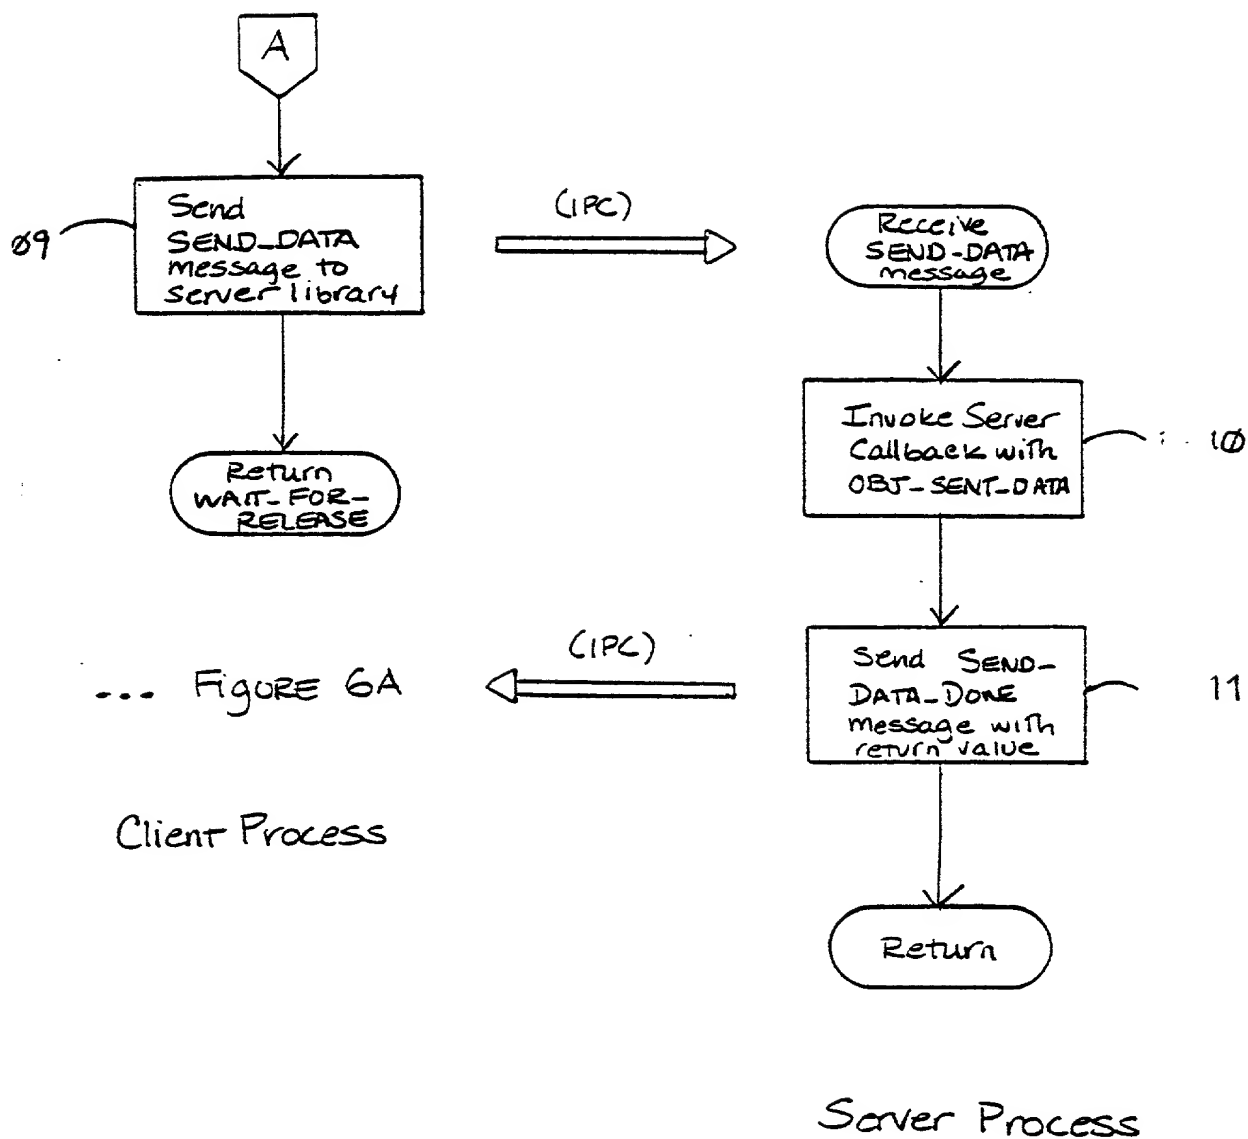


Figure 19 (CONT'D)

DECLARATION AND POWER OF ATTORNEY

As the below-named inventor, I declare that:

My residence, post office address, and citizenship are as stated below under my name.

I believe I am the original, first, and sole inventor of the invention entitled "A METHOD AND SYSTEM FOR REGISTERING DATA FORMATS FOR OBJECTS," which is described and claimed in the specification and claims of patent application Serial No. 07/900,968, which I filed in the United States Patent and Trademark Office on June 17, 1992 and for which a patent is sought.

I have reviewed and understand the contents of the above-identified specification and claims, as amended by any amendment specifically referred to herein (if any).

I acknowledge my duty to disclose information of which I am aware which is material to the examination of this application in accordance with 37 C.F.R. § 1.56(a).

I hereby appoint RICHARD W. SEED, Registration No. 16,557; ROBERT J. BAYNHAM, Registration No. 22,846; EDWARD W. BULCHIS, Registration No. 26,847; GEORGE C. RONDEAU, JR., Registration No. 28,893; DAVID H. DEITS, Registration No. 28,066; WILLIAM O. FERRON, JR., Registration No. 30,633; PAUL T. MEIKLEJOHN, Registration No. 26,569; DAVID J. MAKI, Registration No. 31,392; ROBERT M. STORWICK, Registration No. 30,112; RICHARD G. SHARKEY, Registration No. 32,629; GEORGE B. FOX, Registration No. 31,510; DAVID V. CARLSON, Registration No. 31,153; MAURICE J. PIRIO, Registration No. 33,273; KARL R. HERMANN, Registration No. 33,507; DAVID D. McMASTERS, Registration No. 33,963; ROBERT IANNUCCI, Registration No. 33,514; S. KEVIN PICKENS, Registration No. 34,696; JOSHUA KING, Registration No. 35,570; MICHAEL J. DONOHUE, Registration No. P-35,859; and LORRAINE LINFORD, Registration No. P-35,939, composing the firm of SEED and BERRY, 6300 Columbia Center, Seattle, Washington 98104-7092, my attorneys to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. Please

direct all telephone calls to Maurice J. Pirio at (206) 622-4900 and telecopies to (206) 682-6031.

I further declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further, that these statements were made with the knowledge that the making of willfully false statements and the like is punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and may jeopardize the validity of any patent issuing from this patent application.

AS Williams

Antony S. Williams

Date 8/11/92

Residence : City of Redmond, County of King
State of Washington

Citizenship : United Kingdom

P.O. Address : 22542 N.E. 46th Street
Redmond, Washington 98053